Decision-making beyond "left or right"

A computational study on the neurophysiology behind multiple-choice decision-making and choice reevaluation

Larissa Albantakis

TESI DOCTORAL UPF / 2011

Supervised by

Prof. Gustavo Deco

Department of Information and Communication Technologies

UNIVERSITAT
POMPEU FABRA

Barcelona, September 2011

Acknowledgements

I would hereby like to acknowledge everybody who made this thesis possible. Foremost, of course, is my supervisor Gustavo Deco. During the past four years I profited greatly from his constant motivation, inspiration, and support, both in terms of my PhD thesis, and with respect to all the various summer schools, workshops, and conferences that I could enjoy in the course of my studies.

I am further indebted to a lot of people, both within and outside the university, whose help proved invaluable during this PhD. Direct contributions to the content of this thesis came from Albert Costa and Francesca M. Branzi, my collaborators in the psychophysical experiment (Chapter 5 and 6). Moreover, I am very grateful to:

- The members of my PhD committee, Jaime de la Rocha, Markus Diesmann, and Ruben Moreno-Bote, as well as their substitutes Albert Compte and Salvador Soto-Faraco, for their time and interest.

- Markus Diesmann, for giving me the opportunity to carry out my summer project at the RIKEN BSI under his supervision, and for all the subsequent support regarding my scientific future.

- Christian Lohmann, for his ongoing support through recommendation letters, and the much appreciated discussions about science and life.

Thanks also to my many colleagues for making my PhD a very pleasant time. Especially:

- Joana and Yota, for actually being friends, not colleagues.

- Laura and Marina, for insights into (almost) local issues, and their solidarity in scientific and personal topics.

- Adrian, Mario, and Tristan, for their company, a lot of fun, and advice on all life situations.

- Andres for his never-ending patience with all my minor and major computer issues. What would I have done without you?

- Carolina, Daniel, Elena, and Joana. I could not have wished for better office mates.

- Tim, for having been an ideal coauthor, and for his skills in organizing extracurricular activities.

- Daniel, for his fruitful attempts to expand my cultural horizon.

Abstract

Neurophysiological brain processes during perceptual decision-making have mainly been investigated under the simplified conditions of two-alternative forced-choice (2AFC) tasks. How do established principles of decision-making, obtained from these simple binary tasks, extend to more complex aspects like multiple choice-alternatives and changes of mind? Here, we first address this question theoretically: based on recent experimental findings, we extend a biophysically realistic attractor model of decision-making to account for multiple choice-alternatives and choice reevaluation. Moreover, we complement our computational approach by a psychophysical experiment, exploring how changes of mind depend on the number of choice-alternatives. Our results affirm the general conformance of attractor networks with higher-level neural processes. In particular, we found evidence for the physiological relevance of a so far unregarded bifurcation. Furthermore, our findings suggest an advantage of a pooled multi-neuron representation of choice-alternatives, and a negative correlation between reaction time and changes of mind, possibly regulated by the decision threshold. Finally, we gained testable predictions on neural firing rates during changes of mind and propose future experiments to distinguish nonlinear attractor from linear diffusion models.

Resumen

Los procesos neurofisiológicos que tienen lugar en el cerebro durante la toma de decisiones basadas en fenómenos de percepción han sido investigados, principalmente, en condiciones simplificadas, en particular, de tareas con dos alternativas y elección forzada (2AFC). ¿Cómo podemos extender los principios establecidos sobre la toma de decisiones obtenidas a partir de estas tareas simples y binarias, a aspectos más complejos como decisiones con alternativas múltiples y los cambios de opinión? En esta tesis, en primer lugar, abordamos esta cuestión de manera teórica: a partir de resultados experimentales recientes, extendemos un modelo de toma de decisiones, que es un modelo con atractores realista desde el punto de vista biofísico, con el objetivo de explicar la elección con alternativas múltiples y la reevaluación de la elección. Además, complementamos nuestro enfoque computacional con un experimento psicofísico, explorando cómo los cambios de opinión dependen del número de alternativas. Nuestros resultados refuerzan la tesis de que existe una correspondencia general entre las redes de atractores y los procesos neuronales superiores. En particular, revelan la importancia fisiológica de una bifurcación que hasta ahora ha pasado inadvertida. Además, sugieren la ventaja de representar las alternativas de elección con múltiples neuronas, y la existencia de una correlación negativa entre el tiempo de reacción y los cambios de opinión, posiblemente regulada por el umbral de decisión. Finalmente, proporcionamos predicciones comprobables sobre las tasas de disparo neuronal durante los cambios de la opinión y proponemos experimentos futuros para distinguir los modelos no lineales con atractores de los modelos de difusión lineal.

Preface

How do we make decisions? A comprehensive answer to this multifaceted question cannot be given within the scope of one single scientific discipline. Fields from philosophy to economic sciences aim to shed light on particular aspects of decision-making, approaching the topic from different perspectives: Philosophers first of all seek definitions of the terms "decision" and "choice", but are eventually concerned with the existence, or role, of free will and consciousness in the decision process. Economists take a value-based approach, intending to solve the problem of utility maximization, i.e. the question of what method would lead to the largest reward rate. This is closely related to mathematical decision theory, where Bayesian inference provides a decision rule for optimal choice based on uncertain evidence. By contrast, psychologists strive for phenomenological models describing human and animal decision behavior, which often approximates optimality in the mathematical sense, but can also be (apparently) irrational (Bogacz et al., 2006). Finally, neuroscientists try to gain insights into the actual brain processes during decision-making and to answer how mathematical and psychological concepts are eventually implemented in the brain. With this thesis, we aim to contribute to elucidating this latter problem from a computational point of view: by using mathematical and numerical methods we translate conceptual models of decision-making into biophysically realistic systems. These neural network models are then used to simulate and predict neural activity during the choice process.

Before getting started by reviewing state of the art experimental and modeling approaches on decision-making, it is necessary to define and delimit the kind of decisions we will deal with in the course of this thesis. In all experimental tasks described below, "perceptual" decisions are formed based on evidence in the form of sensory stimuli, which leads to a choice for one of two or more alternative actions. Consequently, "decision" and "choice" here go hand in hand. Generally, however, there is a distinction between the two terms: whereas a choice is the commitment to an alternative, which is indicated by an action made for a certain purpose, a decision refers to the internal deliberation about the alternatives, preceding the choice (Schall, 2001). Interestingly, it is still an open issue to what extent the brain distinguishes decisions and choices in this terminological sense. For perceptual decisions, it is not entirely clear yet, whether or not the decisions are represented on an abstract level, independent of response modality, i.e. independent of the movement that communicates the internal decision. In macaque monkeys, the very same neural populations that are involved in movement preparation show decision-related activity prior to a choice that leads to the respective motor response. For instance, single cell activity in the posterior parietal

cortex depends on whether the choice is indicated by a hand or an eye movement, in an otherwise identical sensory discrimination task (Cui and Andersen, 2007). On the other hand, a recent fMRI study on humans (Heekeren et al., 2006) revealed a candidate brain region for abstract representation of perceptual decisions, the left posterior dorsolateral prefrontal cortex (DLPFC). This invites to the speculation that humans, unlike nonhuman primates, may have evolved a more abstract decision-making circuitry, allowing for higher flexibility between decision and action (Heekeren et al., 2008). Nevertheless, decision-related neural activity recorded in monkeys has been remarkably consistent with neuroimaging studies in humans, despite the differences in techniques (invasive and non-invasive). In both species sensory evidence is accumulated in lower-level sensory regions and compared further down-stream in higher-level brain areas. In sum, insights on neural computations obtained from invasive recordings in monkeys do allow for conclusions on the basics of human decision-making.

This strong similarity across species, as well as our attempt to simulate decision-making and related brain activity using computational models, naturally touches delicate philosophical issues. Yet, we will dismiss this matter here by confining our study to choices solely attributable to the sensory input to - and the assumed properties of - the decision-making unit. Several other factors, such as attention, prior probability of the decision alternatives, or the expected reward, can bias perceptual decisions towards one of the alternatives, but were not explicitly included in any of the experimental or model designs treated here. Instead, this thesis builds on decades of research on the simplest case: basic forced choice tasks with two alternatives. The behavioral and neurophysiological data gained from these simple binary tasks has largely motivated and constrained current models of decision-making. To date there is broad consent that the brain implements (time consuming) perceptual decisions with some sort of integration-to-threshold mechanism, where sensory evidence is accumulated over time until a decision criterion is reached. The details of this decision mechanism, whether it is linear or nonlinear, with more emphasis on early or late evidence, and the way it is actually implemented in the brain, are, however, still open to debate. In order to enhance the understanding of the neural mechanisms underlying perceptual choices, the focus of recent experimental as well as modeling studies shifted to more complex aspects of decision-making, such as multiple choices and changes of mind.

Based on this groundwork (see Chapter 1 and 2), the present thesis means to complement current knowledge of perceptual decision-making by combining theory and experiment in a study of the neural computations behind multiple-choice decision-making, changes of mind, and their combination.

Contents

List of figures

List of tables

1 GENERAL INTRODUCTION

Decision-making generally refers to the process of deliberation in an internal debate about possible choice-alternatives. The criteria that determine the desirability or preference of certain choices are versatile and can differ substantially between individuals. Economic limitations, subjective taste, and previous experience are just a few examples of factors that bias our considerations. Some of them we take into account consciously, others influence us unconsciously.

In order to comprehend the cognitive processes during decision-making, it is important to control causal and subjective factors as much as possible. Therefore, compared to "real-life" decisions, for instance what to choose for lunch, or which shirt to wear, decision-making in psychophysical and neurophysiological experiments is typically reduced to highly simplified conditions.

Here, we are particularly concerned with *sensorimotor* decisions, a special type of "perceptual classification judgments", where one of several motor responses has to be chosen based on some form of sensory evidence in favor or in contra of the possible response alternatives. In the context of sensorimotor choices, the decision process thus corresponds to the translation of perception into action.

Until recently, the basic principles underlying this process have been investigated predominantly in the framework of two-alternative forced-choice (2AFC) tasks. Subjects performing 2AFC tasks must make a choice between two alternatives, which is evaluated based on *choice accuracy* and *reaction speed* (Bogacz et al., 2006). The amount of sensory evidence available to the subject determines the task *difficulty* and influences the behavioral performance. Typically, these parameters are interrelated in the sense that more evidence can be gathered if the decision time is longer and more evidence leads to more accurate choices. Time will improve the probability to make the correct choice, particularly when a perceptual decision has to be made based on noisy, moving, or ambiguous sensory evidence. This time-accuracy relation, termed "speed-accuracy tradeoff", can be investigated in "free response" tasks, where the decision-maker autonomously terminates the evidence accumulation. In this case, some internal decision criterion determines the end of the decision process and elicits the respective motor response.

The notion of gathering noisy sensory evidence over time, until a decision criterion (or "bound") is reached, is termed "accumulation-to-bound" principle. It is incorporated in a class of phenomenological decision-making models summarized as "sequential sampling models", which view the decision process as a *decision variable* evolving in time

until it hits a decision threshold (e.g., Stone, 1960; Laming, 1968; Vickers, 1970; Ratcliff and Smith, 2004). Before decision-related neurophysiological data from 2AFC studies became available, these formal mathematical models aimed at providing mechanistic explanations of the still obscure neural process, restricted by behavioral data from psychophysical 2AFC experiments (Luce, 1991; Ratcliff and Smith, 2004).

A now classic experimental paradigm, designed specifically to test for neural implementations of the "accumulation-to-bound" models, is the random-dot motion (RDM) discrimination task. Strikingly, single-cell recordings from several brain areas along the visuomotor pathway of behaving monkeys indeed revealed potential neural correlates of the theoretical decision variable (reviewed in Opris and Bruce, 2005; Gold and Shadlen, 2007).

Neurophysiological findings further helped to distinguish redundant phenomenological models from biologically plausible models of the neural dynamics underlying decision-making (Ratcliff et al., 2003; Schall, 2003; Smith and Ratcliff, 2004). Moreover, they motivated the development of mathematical models with explicit analogy to neural mechanisms, which aim to account for behavioral and neurophysiological data at once (Wang, 2002; Mazurek et al., 2003; Ditterich, 2006b; Wang, 2008). To date, two models of decision-making proved particularly successful to account for a vast amount of behavioral and neurophysiological data recorded during 2AFC paradigms: the conceptual "drift-diffusion" model (DDM) (Ratcliff and Rouder, 1998), and a biologically-inspired nonlinear attractor model (Wang, 2002).

All in all, behavioral data gained from sensorimotor 2AFC tasks, together with complementary single-neuron recordings, motivated, constrained and advanced models of decision-making during the past decades. This body of experimental and modeling studies forms the groundwork of the present thesis and will be reviewed in Chapter 2.

2AFC tasks, nevertheless, by definition neglect important features of real-life decision-making. First, everyday decisions mostly involve the need to select between not two, but multiple alternatives. Second, decisions are not necessarily absolute but may occasionally be adjusted if we have changed our mind.

The purpose of this thesis is to shed light on the neuronal computations underlying decision-making beyond the limitations of 2AFC tasks. In particular, we ask how established decision models and fundamental concepts like the accumulation-to-bound principle extend to more complex aspects of decision-making, such as multiple choice-alternatives and change of mind. Experimentally, this question has been

addressed just recently, in the context of the RDM task (Churchland et al., 2008; Niwa and Ditterich, 2008; Resulaj et al., 2009).

Using the behavioral and neurophysiological data gathered in these experiments as restrictive evidence, our first objective was to analyze the physiologically realistic attractor model of decision-making in the light of changes of mind and multiple choice-alternatives.

Chapter 3 of this thesis is dedicated to our theoretical account on changes of mind. By making explicit use of the model's nonlinear attractor properties, we were able to replicate psychophysical findings on changes of mind between two alternatives (Resulaj et al., 2009). What is more, as the model is implemented in a biologically-detailed network of spiking neurons, we gained neurophysiological predictions on neural activity during the change process.

In Chapter 4 we will turn to multiple-choice decision-making. There, we propose an extension of the binary attractor model to up to four choice-alternatives. In particular, we increased the number of discrete neural populations, which represent the choice-alternatives in the model. In this way, we could explain all relevant observations from an experimental study that compared decision-related behavior and neural activity of monkeys given two and four choice-alternatives (Churchland et al., 2008). To that end, we analyzed how the network's competition regimes could be brought into accord for different numbers of alternatives. In doing so, it proved beneficial to represent the choice-alternatives by larger neural populations.

While Chapters 3 and 4 treat changes of mind and multiple choices separately, in Chapter 5 we combine these two extensions of classic 2AFC tasks for the first time. In a novel psychophysical experiment and complementary computational analyses, we address the question of how changes of mind depend on the number of choice-alternatives. In short, with more choice alternatives, choice corrections became less likely. Moreover, we found a negative correlation between changes of mind and mean reaction times across participants. An attractor model that combines key features of the model versions applied in Chapter 3 and 4 could explain our behavioral results and even accounted for between subject variations through adaptation of the decision threshold.

Finally, in Chapter 6 we deal with a somewhat puzzling side effect in our participants' behavior during the multiple-choice/changes of mind experiment of Chapter 5, which has previously been omitted for clarity. By means of the attractor model, we could trace this observation back to a visual illusion in the visual motion stimulus. With the additional assumption of an "illusion bias" in the sensory evidence we were able to simulate the participants' behavior in more detail. This further confirms the explanatory power of the physiologically realistic attractor model.

Following our theoretical and experimental findings, Chapter 7 is devoted to a general discussion about the aptitude of our theoretical approach to describe real cortical processes. There, we highlight common implications of our findings with respect to attractor dynamics and discuss the plausibility of our predictions, considering the models' necessarily limited physiological accuracy.

We conclude this dissertation with a summary of our most relevant results and an outlook on future objectives and challenges in Chapter 8.

2 "2AFC" DECISION-MAKING

Binary choices are certainly a simplification of most situations we encounter in our daily lives. Nevertheless, they are representative of many ordinary decision problems, e.g. whether to turn left or right at a crossing, etc. (Bogacz et al., 2006). Throughout the last century, a large collection of psychophysical data has been gathered from 2AFC tasks (e.g., Hill, 1898; Luce, 1991). Based on this groundwork, mechanistic models have been developed to enhance the understanding of still covert neural decision processes by replicating experimentally observed behavior (e.g., Stone, 1960; Laming, 1968; Vickers, 1970; Ratcliff and Smith, 2004).

With recent advances in electrophysiology, in vivo single-cell recordings from primates performing 2AFC decision tasks have become feasible (reviewed in Opris and Bruce, 2005; Gold and Shadlen, 2007). The resulting neurophysiological data further constrained established mechanistic models and motivated sophisticated biologically plausible models of neural processes during decision-making.

In this chapter, we will review experimental and theoretical studies on 2AFC decision-making, which set a precedent to all current attempts to extend our knowledge of perceptual decision-making.

2.1 Experimental Basis

Among all neural systems, sensory and motor circuits are probably the most studied in neuroscience in general. Therefore, it is not surprising, that perceptual decisions, and, in particular, sensorimotor choices, also form the prime subject of study in neuroscientific research on decision-making. Of all senses, primates, including humans, particularly rely on vision to guide their behavior (Opris and Bruce, 2005). In order to react appropriately to a visual scene, we typically need to combine different visual cues, which naturally comprise some uncertainty. This notion gave rise to a now classic visual motion discrimination paradigm, which forms the experimental basis of our work.

2.1.1 Visual motion discrimination task

The so-called "random-dot motion" (RDM) discrimination task is a well-established psychophysical paradigm, designed to study the time course of slow perceptual decision-making (Britten et al., 1992; Roitman and Shadlen, 2002; Palmer et al., 2005; Churchland et al., 2008).

Subjects performing this task have to decide on the net direction of motion in a patch of moving dots (Fig. 2.1). While most dots are moving
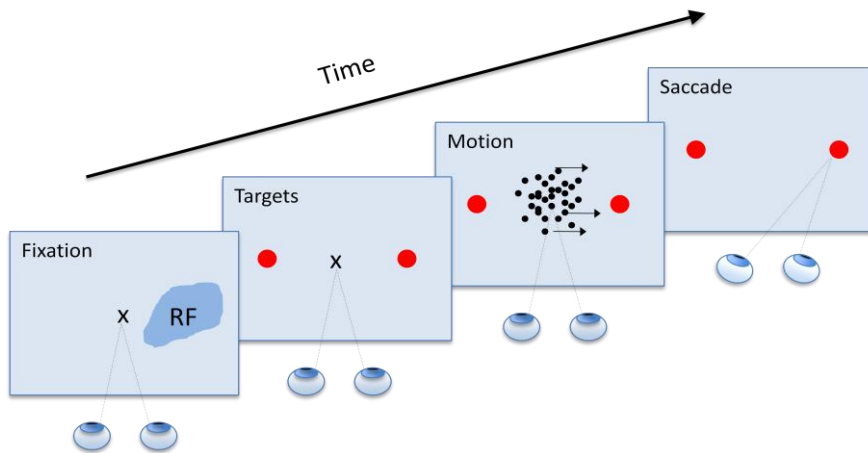
Fig. 2.1 Random-dot motion discrimination task
While the subject is fixating on a central spot, the possible alternatives (R-targets) are indicated. During simultaneous neurophysiological recordings, one of the R-targets (red dots) is located in the response field (RF) of the recorded neuron. After a random delay, a patch of moving dots appears. A certain percentage of these dots are directed towards on of the R-targets, while the others move randomly. The subject reports its choice with a saccade to the corresponding R-target.

randomly, a certain percentage of dots coherently travels in one of several potential directions. The amount of coherent motion thereby controls the quantity of sensory evidence and, thus, the task difficulty. The possible direction alternatives are specified by response targets (R-targets) prior to the onset of the random-dot stimulus. In other words, if more dots are moving towards one of the R-targets, it is easier to detect the correct motion direction. Typically, the subject indicates its choice by a saccadic eye movement to the R-target located in the corresponding direction.

In one version of the RDM task, the "fixed duration", or "interrogation" paradigm, the experimenter determines the end of a trial, instructing the subject to make a choice after a specified time interval. Another more common way to conduct the RDM task, is the "free response" paradigm, where subjects indicate their decision, as soon as they have gathered enough evidence to make a choice. All RDM experiments presented in Chapters 3-6 were carried out in this way. On top of choice *accuracy*, i.e. whether the correct or wrong R-target was chosen, the free response paradigm allows to measure *reaction time* (RT) as a second variable to evaluate behavioral performance. Reaction times are generally determined by the onset of the subjects' motor response. In the case of the RDM task, reaction times are usually long, in the order of several hundred milliseconds, with faster responses to stronger coherent

motion (Roitman and Shadlen, 2002; Palmer et al., 2005; Churchland et al., 2008).

Importantly, the RDM task directly implements the notion of perceptual decision-making as a temporal integration of noisy sensory evidence until a decision criterion is reached, the so-called "accumulation-to-bound" principle. Due to the randomly moving dots, the momentary amount of coherent motion is subject to stochastic fluctuations. The correct direction of motion can thus be inferred more reliably the longer the motion stimulus is viewed. Consequently, the reaction time and accuracy of a decision are not independent of each other (Palmer et al., 2005). This relation is generally known as the "speed-accuracy tradeoff" (SAT).

Recently, the RDM task has been modified independently in several ways: Churchland et al. (2008) and Niwa and Ditterich (2008) extended the RDM task from binary to multiple choices. In particular, Churchland et al. (2008) compared monkeys' behavioral and neurophysiological responses between a 2- and 4-alternative version of the RDM discrimination task. Reaction times and error rates for four alternatives were found to be longer and higher, respectively, consistent with earlier studies on multiple-choice decisions (Hick, 1952). Notably, the experiments of Churchland et al. (2008) provided the first electro-physiological data of a 4-alternative decision task. We will present their results in detail in Chapter 4, as they provide the experimental comparison for our multiple-choice modeling efforts.

Niwa and Ditterich (2008) tested human participants on a 3-alternative version of the RDM task with a multicomponent RDM stimulus, which was comprised of up to three coherent motion components instead of just one direction of coherent motion. This additional feature allowed them to control the amount of sensory evidence for all three choice-alternatives, creating situations with identical choice performance but different reaction times. For example they found that for identical motion strength in all three directions responses were faster than without coherent motion, although the net evidence was the same, namely zero, leading to chance level performance. The multicomponent experiment might prospectively help to distinguish modeling approaches, especially in combination with neurophysiological recordings (see 2.3 and 3.4.1).

Finally, Resulaj et al. (2009) modified the RDM task in order to study "changes of mind" due to further processing of available information after an initial decision. Instead of responding with a saccade to the chosen target, as in the standard RDM task, human participants were asked to move a handle to a left or right target. The reason is that only with a continuous movement changes of mind could directly be observed in the movement trajectory. A saccade, on the other hand, is a rather ballistic

movement. This experiment will be presented in detail in Chapter 3 as a basis for our theoretical account of changes of mind.

2.1.2 Neural correlates of decision-making

To identify possible neural correlates of the accumulation-to-bound concept, the psychophysical RDM paradigm was combined with simultaneous recordings of decision-related single-cell activity. Taking visual motion discrimination as a prototype for perceptual decision-making has the advantage that anatomical and functional properties of the visuomotor pathway are particularly well determined. During the last two decades, several studies successively targeted brain areas along the cognitive link between visual sensation and saccadic movement in search of neurons that might encode a decision variable (reviewed in Schall, 2003; Smith and Ratcliff, 2004; Opris and Bruce, 2005). Fig. 2.2 shows a simplified schematic diagram of the cortical and subcortical neural areas involved in the processing of visual information (colored in blue) and the execution of saccadic eye movements (colored in orange).

Generally, visual information arriving in the primary visual cortex (V1) is further processed via two specialized pathways: first, the *ventral stream* associated with forms and colors, involved in "what" tasks like object recognition, and, second, the *dorsal stream*, which is processing "where" information necessary for motion detection. More precisely however, the visual areas form a complex network and also the two main processing pathways are strongly interconnected. As motion information is the only relevant stimulus feature in the RDM task described above, in the following we will focus on characterizing the cortical areas along the dorsal stream.

The first cortical region down-stream of the visual cortex, which holds information about the direction of motion, is the middle temporal area (MT) (Fig. 2.3A, inset in C). Area MT encodes the absolute amount of visual motion. Neurons in MT have direction-sensitive tuning curves with one preferred direction of motion. Presented with a RDM stimulus, MT neurons will fire more, the higher the coherent motion in their preferred direction and less, for coherent motion in their null direction. In many MT neurons this relation between firing rate and motion coherence is approximately linear (Britten et al., 1992, 1993).

The absolute motion information present in area MT is still insufficient to explain a subject's motor response. Instead, neural activity from area MT might act as the source of evidence upon which further down-stream areas base their choice. This view was further confirmed by electrical stimulation of MT neurons from monkeys performing the RDM task. The monkeys' choices were biased towards the preferred direction of the stimulated MT neurons (Salzman et al., 1992; Ditterich et al., 2003).
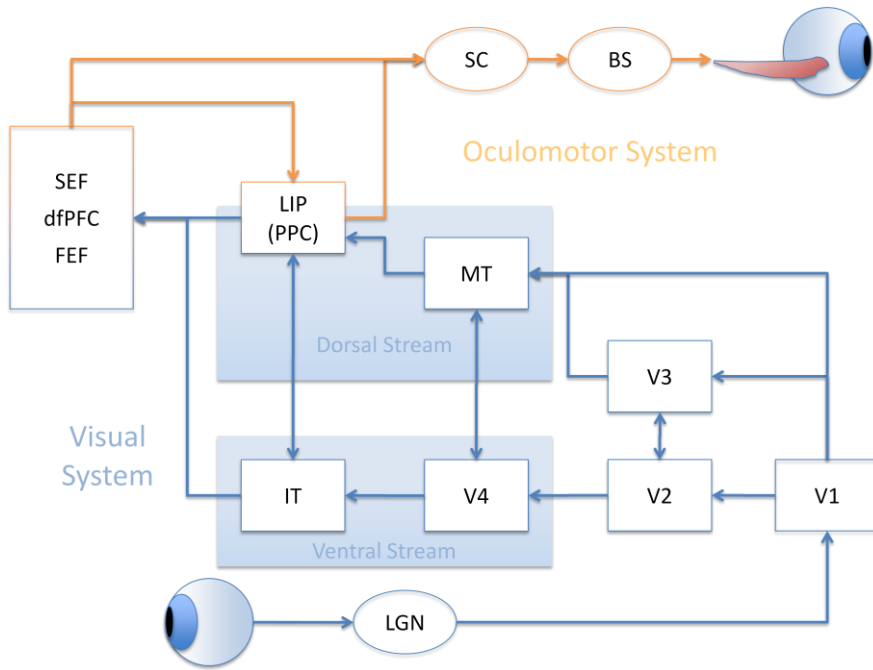
Fig. 2.2 Neural circuitry engaged in visual discrimination tasks

Visual signals from the retina arrive in the primary visual cortex (V1) through the lateral geniculate nucleus (LGN). V1, V2, V3, and V4 are primary, secondary, tertiary, and quaternary visual areas. "What" information about visual stimuli, like form and color, is further processed via the ventral stream, passing through V4 and the inferior temporal (IT) cortex. "Where" information necessary for motion detection is sent from V1 over the middle temporal area (MT) to the lateral intra-parietal area (LIP) located in the posterior parietal cortex (PPC). Information from the two pathways is combined in the prefrontal cortex (PFC), in particular in the frontal eye field (FEF), the supplementary eye field (SEF) and the dorsolateral prefrontal cortex (dlPFC). The command to execute a saccade from PFC or LIP is passed through the superior colliculus (SC) to the brainstem (BS), which activates the ocular muscles. Adapted from (Opris and Bruce, 2005).

Stimulation thereby shifted the *psychometric function*, the dependence of accuracy on motion coherence. Strikingly, it also affected reaction time: choices in the neurons' preferred direction were speeded up by electrical stimulation, while choices in their null direction were slowed down. Stimulated neural activity in MT thus influences decision behavior in the same way as additional visual evidence would.

   Next in line along the dorsal stream and directly innervated by area MT, lies the lateral intraparietal area (LIP) within the posterior parietal cortex (PPC). The activity of most neurons in PPC is both sensory- and

motor-related and can be associated with the formation of intentional motor plans (reviewed in Andersen, 1995). In particular, LIP neurons fire prior to a saccade directed into their "response field" (RF), but also respond to static visual stimuli located in their RF.

Based on these initial findings, Shadlen and colleagues first combined the RDM paradigm with single-cell recordings from area LIP of behaving monkeys (Shadlen and Newsome, 2001; Roitman and Shadlen, 2002). In their experiments, one of the visual response targets (R-targets), which indicate the possible motion directions, was always placed in the RF of the recorded LIP neuron. Thereby, single-cell activity of LIP neurons has been found to increase gradually during motion viewing if the monkey subsequently chose the R-target inside the response field, and to decline comparably if the opposite R-target was chosen. Moreover, the slope of this activity build-up depends on task difficulty, i.e. the amount of motion coherence (Fig. 2.3B,C)[1]. LIP neurons thus seem to accumulate, or integrate incoming information about motion direction according to choice behavior. In contrast to more up-stream areas, LIP neurons hence show all signs of a potential neural decision variable.

Besides, strikingly consistent with the accumulation-to-bound principle, LIP activity suggests a fixed firing rate threshold, as it reaches a uniform level about 40-80 ms prior to the saccade, independent of response time, difficulty, and even the number of alternatives (Fig. 2.3D) (Roitman and Shadlen, 2002; Churchland et al., 2008). How this decision threshold might be regulated or read out in the brain is still largely unknown. Recent theories about possible neural substrates of the decision threshold involve cortico-collicular and cortico-basal ganglia circuits (Lo and Wang, 2006; Bogacz and Gurney, 2007).

Decision-related build-up activity during the RDM paradigm was also found downstream of LIP, in the dorsolateral prefrontal cortex (dlPFC) and the superior colliculus (SC) (Horwitz and Newsome, 1999; Kim and Shadlen, 1999). Interestingly, neurons that exhibit such ramping activity characteristically also show persistent neural firing in delayed memory or decision tasks (Gnadt and Andersen, 1988; Schall, 2001).

Whether any or all of these areas (LIP, dlPFC and SC) take an active role in the sensorimotor decision process is still unknown. Some cortical regions might simply reflect the integration of evidence performed in another part of the brain. Furthermore, response preparation might be an alternative explanation for ramping activity, especially for areas further down-stream the oculomotor pathway. In that sense area LIP takes a

---

[1] Note that with "ramping activity" we refer to the average activity across trials. Whether single trial activity builds up gradually, or changes rather abruptly is hard to determine. Future multiunit recordings might provide information on how the population average across neurons in a single trial compares to the trial average of a single neuron.

special role, because it is the first region in the visuomotor chain that exhibits ramping activity, and, in contrast to other regions, a large proportion of its neurons actually show the gradual activity build-up and spatially selective persistent activity (Shadlen and Newsome, 2001; Roitman and Shadlen, 2002; Churchland et al., 2008).
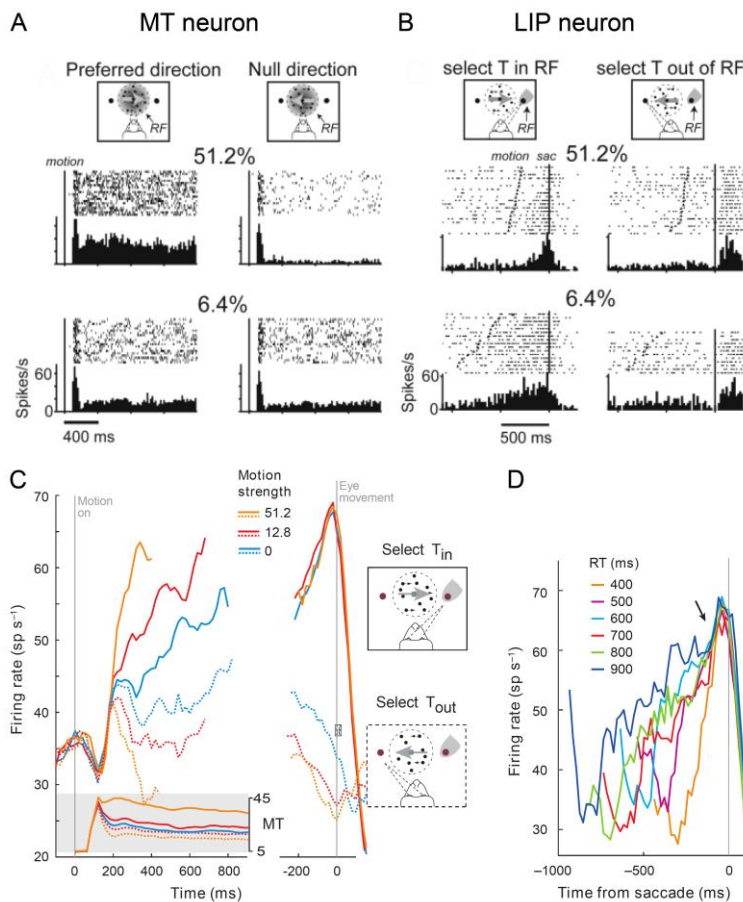
Fig. 2.3 Activity of MT and LIP neurons during the RDM task.
(A) Response of direction selective MT neuron aligned to motion onset (Britten et al., 1992). The RDM stimulus was placed in the receptive field of the neuron. (B) Response of LIP neuron aligned to saccade onset (Roitman and Shadlen, 2002). One choice target was always placed in the response field of the neuron. (A,B) adapted from (Mazurek et al., 2003). (C) Responses of 54 LIP neurons shown for three levels of difficulty and grouped by direction of choice, as indicated. Shaded inset shows average responses from direction selective MT neurons. (D) Responses grouped by RT (only T_{in}). All trials reach a stereotyped firing rate ~70 ms before saccade initiation. (C,D) adapted from (Gold and Shadlen, 2007).

A major component of our approach is to replicate and predict decision-related neural activity during extended versions of the RDM task. With a physiologically plausible spiking-neuron model of decision-making we aim to simulate the entire time course of LIP activity during the different RDM trial phases, as described in (Shadlen and Newsome, 2001; Roitman and Shadlen, 2002; Churchland et al., 2008; Kiani et al., 2008). In the following list we recapitulate the essential points:

- LIP receives direct input from direction selective MT neurons, which fire monotonically as a function of motion coherence.

- LIP neurons strongly respond to the appearance of the visual R-target in their response field.

- With the onset of the RDM stimulus, a "dip" in firing rate occurs, possibly due to divided attention or a top-down reset of activity.

- Starting with a latency of ~190 ms after RDM onset, LIP firing rates gradually rise or decline according to choice-behavior and motion coherence.

- A stereotyped level of activity is reached ~40-80 ms before saccade onset.

- LIP neurons show persistent activity in delayed decision tasks.

With respect to changes of mind, it is worth emphasizing that the neural activity described above and in the following chapters is not confined solely to LIP neurons. As mentioned in the last section, in the RDM experiments on changes of mind (Chapter 3 and 5), participants performed arm movements instead of saccades to indicate their choice. Yet, LIP neurons are mostly associated with saccadic motor responses. Nevertheless, other areas in PPC, especially the parietal reach region (PRR) involved in the preparation of arm movements, share the neural characteristics listed above. More precisely, neurons in PRR show sustained activity during delayed reach to target tasks and also exhibit huge responses to the appearance of a visual reach target in their response field, very similar in size and time course to LIP neurons for saccades in the same paradigm (Snyder et al., 1997; Cui and Andersen, 2007; Andersen and Cui, 2009). Besides, Cui and Andersen (2007) reported that, although generally LIP seems to respond more to eye and PRR more to arm movements, if monkeys are free to choose the motor response, a substantial number of LIP neurons responded preferably to arm movements for instructed motor responses. In sum, the assumptions and predictions on neural activity presented in this section apply generally to both LIP and PRR.

As a final note before turning to the theoretical basis of our work, it has to be clarified that, although in this thesis we focus on the RDM paradigm and visual motion discrimination to represent sensorimotor decision-making, the same general principles apply to other sensory modalities and 2AFC paradigms. Aside from the RDM literature, there is a second, immensely rich body of work in the field of perceptual decision-making, gathered by Romo and colleagues, which addresses sequential decision-making in a tactile discrimination paradigm (reviewed in Romo and Salinas, 2003; Hernandez et al., 2010). There, contrary to motion discrimination, the evidence for or against the two choice-alternatives is not presented at the same time, in parallel, but one after the other with a time delay of several seconds between the two tactile stimuli. The decision process in this sequential setting consequently involves keeping the first stimulus in working memory and deciding based on stored and ongoing sensory information. Apart from this additional complication, the processes underlying the final choice can be described by the similar theoretical concepts as for parallel decisions, which will be reviewed in the following.

2.2 Theoretical Basis

Corresponding to the characteristics of the RDM paradigm, 2AFC models typically make the fundamental assumptions that noisy evidence, subject to random fluctuations, is integrated over time for each alternative, until sufficient evidence has accumulated to make a decision (Bogacz et al., 2006). In the following, we will review the most common models of 2AFC decision-making and their theoretical origins. Thereby, we will start with basic, linear, conceptual models, which successfully capture decision-behavior (2.2.1), followed by attempts to implement these models in a physiologically plausible way (2.2.2). Finally, we will turn to nonlinear attractor models and describe a biophysically inspired implementation of an attractor model with spiking neurons (Wang, 2002), which forms the basis of the models presented in this thesis (2.2.3). Our objective is to provide an intuitive overview. Consequently, we restrict our formal presentation to basic equations and characteristic model features and refer to the original publications for detailed theoretical analysis.

2.2.1 Sequential-sampling models

Present conceptual models of decision behavior considering noisy evidence build on "signal detection theory" (SDT), developed to describe categorical choices under uncertainty (Tanner and Swets, 1954; Green and Swets, 1966). SDT typically assumes fixed, short stimulus times that are

out of the subject's control. The class of models summarized as "sequential sampling models" forms the logical extension of SDT to temporally stretched streams of (noisy) data (Wald, 1947; Stone, 1960). In addition to the probability of correct responses, these models give predictions on subjects' reaction times in "free response" 2AFC paradigms. To form a decision, evidence for each of the two alternatives is integrated over time. Whether an independent integration for each alternative (e.g. race model), or an integration of the difference in evidence (e.g. drift diffusion model) gives a better account of experimental 2AFC data, is, however, still open to debate, although the latter seems to fit a wider set of experimental observations (Ratcliff et al., 2003; Ratcliff and Smith, 2004; Bogacz et al., 2006).

a) Signal detection theory and the SPRT

In simple perceptual 2AFC tasks, subjects are often faced with problems such as: "Has a dim light been flashed or not?" Or: "Which of two similar images has been presented?" Signal detection theory (SDT) provides a prescript for these kind of decisions, where one of two hypotheses has to be chosen on the basis of a single observation in the presence of uncertainty, or noise (Gold and Shadlen, 2007). If the sensory observation is informative about the hypotheses, it provides "evidence" favoring one alternative. We will generally refer to information which is indicative of a choice as evidence e. The two hypotheses H_1 and H_2 stand for the two choice-alternatives. The conditional probability $p(e|H_1)$ denotes the probability of observing evidence e if H_1 is true.

Depending on the signal-to-noise ratio (μ/σ) and the similarity of the hypotheses $(\mu_1-\mu_2)$, the probability density functions (PDFs) of the two alternatives overlap to some degree (Fig. 2.4). The smaller the signal-to-noise ration, the higher the overlap of the PDF. Likewise, the more distinguishable the stimuli, the smaller the overlap. In the case of sensory stimuli the PDFs are often assumed to be normally distributed with means $\mu_1 \neq \mu_2$ and standard deviations $\sigma_1 = \sigma_2$.

The *a posteriori* probability $p(H_1|e)$ that hypothesis H_1 is true given the evidence e can be determined according to Bayes' theorem from the conditional probability $p(e|H_1)$, the *prior*, or *a priori* probability of the hypothesis $p(H_1)$, and the total probability of the evidence $p(e)$:

$$p(H_1|e) = \frac{p(e|H_1)p(H_1)}{p(e)} .$$
(2.1)

The prior $p(H_1)$ thereby denotes the probability that H_1 is true before any evidence has been obtained. If equal priors are assumed for both alternatives, H_1 is more likely to be correct than H_2 if the "likelihood ratio" $LR(e) = p(e|H_1)/ p(e|H_2)$ is larger than 1.
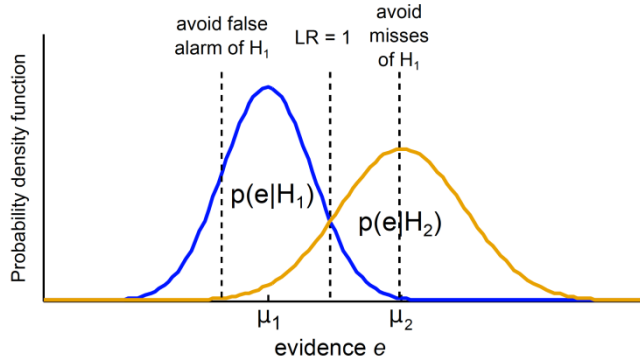
Fig. 2.4 Signal detection theory in 2AFC tasks.
Because of uncertainty the PDFs of the two alternative hypotheses overlap. A choice is made depending on the desired level of accuracy for one of the alternatives. Comparing the likelihood ratio LR to 1 minimizes the total number of errors.

Choosing H_1 if $LR > 1$ is the optimal strategy, in the sense that it provides the lowest overall error rate. In the case of equal rewards or costs, it also indicates the optimal choice in terms of the highest reward.

For some decisions, however, the consequences of a false alarm, for example, are negligible compared to missing a signal. Because of the noise, mistakes are inevitably. Still, the kind of errors, i.e. *false alarms* or *misses*, can be adjusted by the decision criterion (Fig. 2.4). Generally, the desired level of accuracy for one of the alternatives determines the decision threshold, or bound, *B*. For unequal prior probabilities, but identical rewards, H_1 should be chosen if $LR(e) > B = p(H_2)/p(H_1)$ (Green and Swets, 1966). In sensorimotor tasks, unequal priors arise for instance if one stimulus is presented more often than the other.

If not just one, but multiple pieces of evidence $e_1...e_N$ are available over time, as for instance in the random-dot motion (RDM) task, the likelihood ratio has to be updated with each new sample of evidence. With the simplifying assumption that all evidence samples $e_1...e_N$ are independent, the likelihood ratio extends to

$$LR(e) = \frac{p(H_1|e_1\ldots e_N)}{p(H_2|e_1\ldots e_N)} = \prod_{n=1}^{N} \frac{p(e_n|H_1)}{p(e_n|H_2)}.$$

(2.2)

A decision bound $B = 1$ again minimizes the error rate, as it determines the most likely hypothesis (Neyman and Pearson, 1933). From the perspective of 2AFC problems, Eq. 2.2 applies to the "interrogation" paradigm, where the decision is based on a fixed sample of evidence.

In the "free response" paradigm, where the decision-maker is allowed to control the decision time autonomously, she or he is faced with the additional problem *when* to end the evidence accumulation. Accordingly, optimality in free response tasks is often assessed as the strategy that yields the shortest expected reaction time (RT) for a given error rate.

The sequential probability ratio test (SPRT) provides a solution to this specific optimality problem (Wald, 1947). Here, the momentary likelihood ratio *LR(e)* is again calculated as in Eq. 2.2, but instead of one, there are now two decision bounds B_1 and B_2 and the sampling process continues as long as

$$B_2 < LR(e) < B_1, \text{with } B_2 < B_1.$$
(2.3)

In other words, if B_1 is crossed, alternative 1 is selected, if B_2 is crossed, alternative 2 is selected, and while the evidence for both alternatives is insufficient, meaning below a certain level of significance, the decision process continues. Interestingly, a decision rule equivalent to Eq. 2.3 can be obtained using any quantity that is monotonically related to the LR if B is scaled appropriately (Green and Swets, 1966). Hence, by taking the logarithm of Eq. 2.2 and 2.3, the decision process can be written as a simple addition:

$$\log(B_2) < \log(LR(e)) = \sum_{n=1}^{N} \log \frac{p(e_n|H_1)}{p(e_n|H_2)} < \log(B_1).$$
(2.4)

Moreover, the temporal evolution of the log-likelihood ratio (logLR) can be described as a discrete decision variable V, starting at $V_0 = 0$, which is subsequently updated at each time step, according to:

$$V^n = V^{n-1} + \log \frac{p(e_n|H_1)}{p(e_n|H_2)}.$$
(2.5)

Using the logLR to express the SPRT has the further advantage, that evidence in favor of H_1 intuitively adds to V with a positive value, while evidence supporting H_2 contributes negatively. In that sense, the trajectory of the decision variable *V(t)* for noisy evidence is analogous to a one-dimensional "random walk" bounded by a positive and negative threshold.

In the limit of infinitesimally small time steps, equivalent to continuous sampling, the discrete SPRT/random walk model converges to the drift diffusion model (DDM) described in the next section. For a more detailed theoretical description of optimality, also in the case of unequal priors, and the continuum limit of the SPRT please refer to (Bogacz et al., 2006).

Before we turn to the DDM, we briefly discuss how the theory presented above might relate to real neural computations during decision-making and the RDM task in particular. As we have seen in Section 2.1.2, decision-related neural activity in area LIP is consistent with the notion of

an "accumulation-to-bound", while area MT encodes the absolute amount of visual motion present in the RDM stimulus and might consequently provide decisive evidence to LIP. Could LIP activity actually correspond to a neural decision variable in the mathematical sense of the SPRT?

As the brain hardly stores the complete distribution of possible neural responses to every encountered stimulus, it probably has no access to the PDFs of the neural populations, which would be necessary to infer the likelihood ratio LR (Gold and Shadlen, 2001; but see Ma et al., 2006).

However, motivated by the apparent analogy between the trajectory of V and LIP firing rates, Gold and Shadlen (2001, 2002, 2007) argued that a quantity approximating the logLR could indeed be computed in the brain. More precisely, knowledge about the PDFs would not be explicitly necessary to implement a decision rule approximating the optimal SPRT, if output firing rates of two antagonistic sensory neurons or neural populations were used as evidence. One example would be the responses I_1 and I_2 of two populations of MT neurons, one selective for rightward, the other for leftward motion, which respond to their preferred and null direction with the mean firing rates $\mu_1 > \mu_2$, and roughly equal variance σ. In that case, the optimal logLR decision rule will depend only on the firing rate difference I_1-I_2, apart from a scaling factor.

This largely hold true for a variety of possible PDFs (Gold and Shadlen, 2001). In particular:

$$\log LR_{left,right} = \frac{\mu_1 - \mu_2}{\sigma^2}(I_1 - I_2)$$

(2.6)

if I_1 and I_2 are sampled from normal distributions, which is a plausible assumption for the average firing rate of a neural population. Yet, responses of single neurons might better be described by a Poisson distribution. In that case:

$$\log LR_{left,right} = \log(\mu_1/\mu_2)(I_1 - I_2)$$

(2.7)

Knowing the sign of the difference I_1-I_2 in MT activity would hence be sufficient for downstream areas like LIP to elicit a left or right saccade according to an SPRT optimal rule.

Furthermore, a study by Platt and Glimcher (1999) revealed that both, the prior probability of getting a reward, and the expected magnitude of the reward could modulate LIP activity, affirming the suggestion that LIP activity might be a neural correlate of the decision-variable V (Eq. 2.5).

As a final note on the SPRT, the argument of Gold and Shadlen (2001, 2002, 2007) can also be extended to multiple alternatives, or neural populations, which results in a comparison between the neural population with the highest rate and the average rate of the other populations ("max-vs-average test") (Ditterich, 2010). However, contrary to the 2-alternative case, the resulting statistical test is not optimal. Interestingly, the optimal

algorithm for decisions between more than two alternatives is still unknown (McMillen and Holmes, 2006). The multihypothesis (M)SPRT was shown to approximate optimality for small error rates (Dragalin, 1999). For moderate error rates, the physiologically plausible max-vs-average test performs almost as well as the MSPRT (Ditterich, 2010).

b) The drift-diffusion model (DDM)

The continuum limit of the SPRT represents the most basic form of the DDM. A continuous decision variable $v(t)$ is accumulating the evidence difference between the two choice-alternatives, or hypotheses (Stone, 1960; Laming, 1968; Ratcliff, 1978). In the unbiased case with equal prior probabilities, $v(t)$ is integrated over time according to

$$dv(t) = \mu dt + \sigma dW, \quad v(0) = 0,$$
(2.8)

with symmetric decision bounds $b_1 = -b_2$, and the accumulation time interval dt (assumed to be very small). Eq. 2.8 is the continuous extension of Eq. 2.5. The right side of Eq. 2.8 denotes the new noisy evidence obtained during dt. It is composed of a constant drift μdt, with drift rate μ, and the diffusion term[2] σdW, which represents white noise drawn from a Gaussian distribution with mean 0 and variance $\sigma^2 dt$. The correct alternative is determined by μ, which, in the case of the RDM task, can be interpreted as the amount of coherent motion. Using the terminology of the SPRT, if $\mu > 0$, H_1 is correct; if $\mu < 0$, H_2 is correct. Which alternative is eventually chosen by the DDM, however, is also subject to noise, depending on σ, and the height of the decision bounds b_1 and b_2. Still, the DDM, as a continuous implementation of the SPRT, solves 2AFC problems optimally: it will on average return a decision in the shortest possible time for a specified level of accuracy (Bogacz et al., 2006).

Solutions of Eq. 2.8 are normally distributed with probability density $p(v,t) = N(\mu t, \sigma \sqrt{t})$, if the decision bounds are ignored (Gardiner, 1985). Consequently, the variance across trials of the temporal evolution of v increases linearly with t. This can be appreciated in Fig. 2.5, where several example trials of v and their variance are displayed.

Due to the threshold nonlinearity of the decision bounds, the reaction time distributions of correct and error trials are typically left-skewed, with equal means (Fig. 2.5B). Treated as a so called "first-passage time problem" (Ratcliff, 1978; Bogacz et al., 2006), error rates (ER) and mean RTs of the basic DDM can be expressed as:

$$ER = \frac{1}{1 + \exp^{2\mu b/\sigma^2}},$$
(2.9)

---

[2] dW is proportional to $N(0,1) \cdot \sqrt{dt}$, as the variance of uncorrelated stochastic random variables is additive with successive time steps, which leads to a square-root behavior for the standard deviation (Usher and McClelland, 2001).

$$mRT = \frac{b}{\mu} \tanh\left(\frac{\mu b}{\sigma^2}\right) + t_{ND}$$

<div align="right">(2.10)</div>

where t_{ND} denotes the "non-decision" time (e.g. the time related to sensory and motor processes which add to RT).

Contrary to the theoretical predictions of the basic DDM, error responses in 2AFC tasks can have significantly different mean RTs than correct responses, depending on experimental specifications, e.g. stressing accuracy or speed, or the difficulty of a condition, (Luce, 1991; Ratcliff and Smith, 2004; Pleskac and Busemeyer, 2010). A more general, extended version of the DDM includes trial-to-trial variability in the drift rate and the starting point (Ratcliff and Rouder, 1998; Ratcliff and McKoon, 2008). A normally distributed drift rate with mean μ and standard deviation s_μ leads to longer RTs on error trials, as errors will occur more often in trials where the drift μ is small. Drawing the starting point $v(0)$ from a uniform distribution ranging from $-s_v$ to s_v produces on average shorter error RTs, because errors are more likely for a bias
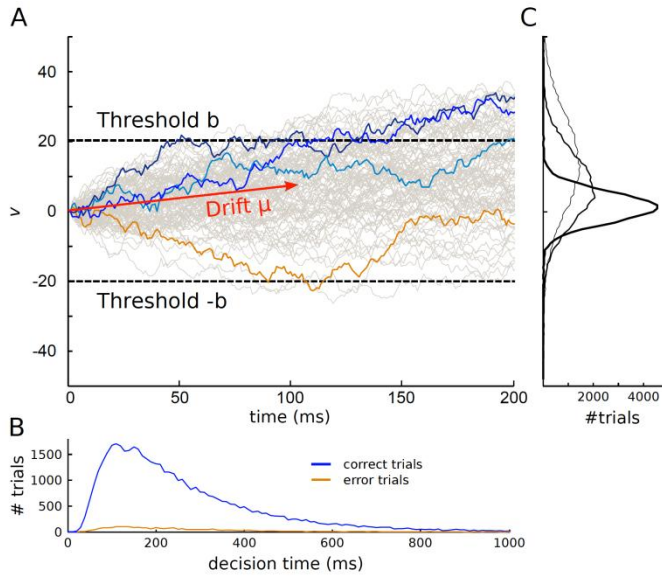
Fig. 2.5 Basic drift diffusion model.
(A) 100 example traces of the time evolution of v(t). Three correct trials were labeled in blue, one error trial in orange. (B) Left-skewed RT Histograms of correct and error choices from 50,000 trials. (C) The variance of v increases with time. The distribution of v for 50,000 trials is given for t =20 ms, 100 ms and 200 ms (bold to narrow lines). Model parameters: μ = 0.07, σ = 1, dt = 1ms, b=20.

towards bound *b2* and hence reach the threshold faster. Physiologically, this variability can be explained by trial-to-trial differences in attention, premature sampling, or other variable perceptual biases.

The extended DDM version is hardly tractable analytically. RT distributions and ERs can only be obtained numerically. Typically, the DDM is fitted to a particular set of behavioral data by minimizing the deviation between a simulated data set and the experimental data set (Vandekerckhove and Tuerlinckx, 2007).

A simplified deterministic version of the DDM has also been proposed by Reddi and Carpenter (2000). Their "LATER" model produces variability in reaction times by varying the drift rate across trials, without any within-trial noise. For the RDM task, where the stimulus itself is explicitly stochastic, this might not seem a plausible model. Within-trials noise, however, is no essential property to fit behavioral data of 2AFC tasks (Brown and Heathcote, 2008).

Yet, the DDM has the advantage that reaction times and accuracy are directly related over the decision threshold. In experiments, more pressure for speed typically leads to faster RTs and lower accuracy. This negative correlation is known as the "speed-accuracy tradeoff" (SAT). By adjusting the decision bounds, the DDM can reproduce the negative SAT correlation (Palmer et al., 2005). Absent noise, as for the LATER model, the threshold has no effect on accuracy.

The reverse conclusion of the SAT is that perfect accuracy could be achieved with unlimited processing time. Accordingly, the DDM implements perfect integration of the evidence difference in the sense that no information is "forgotten" or overly emphasized. Nonetheless, participants' accuracy as a function of RT often reaches an asymptote, especially in difficult tasks and the assumption of perfect integration might not holds true for neural systems. This discrepancies can be solved by introducing a "leakage" term λv to the drift rate in Eq. 2.8. Information loss over time is modeled according to:

$$dv(t) = (\lambda v + \mu)dt + \sigma dW, \quad v(0) = 0, \tag{2.11}$$

with $\lambda < 0$, corresponding to a stable Ornstein-Uhlenbeck (O-U) process (Busemeyer and Townsend, 1993). This can be pictured by a diffusion of v in a curved instead of a flat landscape, where v approaches a stable fixed point $v^* = -\mu/\lambda$, where $dv = 0$. In the opposite case of $\lambda > 0$ both mean and variance of v grow exponentially, as v is repelled from the now unstable fixed point (unstable O-U process).

c) The race model

While in the DDM a single integrator accumulates the evidence difference between two alternatives, in the race model (Vickers, 1970, 1979) separate integrators v_1, v_2 are used for each alternative:
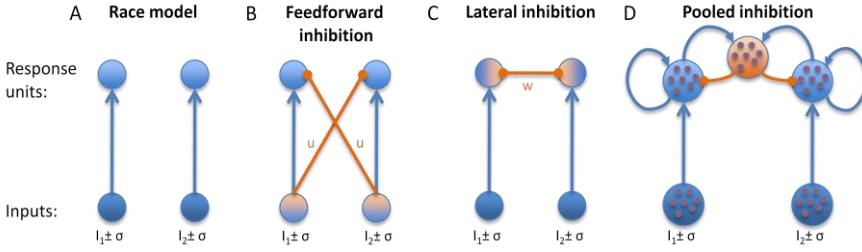
Fig. 2.6 2AFC decision models with two integrators.
Blue color denotes excitatory connections and populations, orange inhibitory. Adapted from (Bogacz et al., 2006).

$$dv_1 = I_1 dt + \sigma dW_1$$
$$, \quad v_1(0) = v_2(0) = 0,$$
$$dv_2 = I_2 dt + \sigma dW_2$$

$$(2.12)$$

I_1 and I_2 denote the average incoming evidence, respectively. As for the DDM, white noise is sampled from a normal distribution, $N(0, \sigma^2 dt)$. In the free-response mode, a decision is made as soon as one of the two integrators exceeds a threshold B[3]. The two integrators thus perform a "race to threshold" against each other.

The race model is not equivalent to the DDM and thus not optimal (Bogacz et al., 2006). Nevertheless, other than the DDM, it can easily be extended to multiple-choice decisions, simply by adding more integrators.

Moreover, in the race model, v can be interpreted as the population activity of two neural pools, receiving inputs from distinct populations of up-stream sensory neurons. For the DDM, however, it remains unclear how the difference in evidence might be computed physiologically.

This problem has been addressed in subsequent "connectionist" models of 2AFC decision-making (Fig. 2.6). These abstract neural network models implement the diffusion process with inhibition between two integrators and will be reviewed in the following section.

2.2.2 Biologically-motivated rate models

As we have seen, the DDM is an intuitively appealing model of 2AFC decision-making and, moreover, achieves optimality according to the SPRT. But, is there a way to implement this drift-diffusion concept in a physiologically plausible manner?

---

[3] We again assume equal prior probabilities. Therefore both integrators have the same decision bound B.

Several models have been proposed, which effectively compute evidence subtraction with different inhibitory mechanisms (Fig. 2.6). Although these models all exhibit two separate integrators just like the race model, dynamically they are more closely related to the DDM (Bogacz et al., 2006).

a) Feedforward inhibition (FFI)

Mazurek et al. (2003) proposed a feedforward inhibition (FFI) model directly motivated by neural activity from MT and LIP during the RDM motion task (Britten et al., 1993; Shadlen and Newsome, 2001). The model instantiates the hypothesis that two populations of movement-sensitive MT neurons provide evidence in favor or against the opposed motion directions, which is then integrated in area LIP (Fig. 2.6B). A decision is made, when the pooled activity reaches a threshold. The two LIP populations thus correspond to the response units, which receive excitatory inputs from one population of MT neurons and inhibitory inputs from the other, according to:

$$dv_1 = I_1 dt + \sigma dW_1 - u(I_2 dt + \sigma dW_2)$$
$$dv_2 = I_2 dt + \sigma dW_2 - u(I_1 dt + \sigma dW_1)$$

$$, \quad v_1(0) = v_2(0) = 0,$$

$$(2.13)$$

where u denotes the inhibitory FF connection weight. In this simple version of the FFI model, integration is perfect without leak.

Just as the basic DDM model, the FFI model cannot account for the slower error RTs found experimentally in the RDM task without further extensions. Ditterich (2006a; 2006b) consequently suggested that a time-variant version of the FFI, including for example a within-trial increase of the input gain, could account quite well quantitatively for both behavioral and neural data of Roitman and Shadlen (2002). An additional leak term further improved the fit to the neural data.

Besides, Niwa and Ditterich (2008) successfully applied a 3-alternative version of the FFI model to their RDM experiment with three possible directions of motion (see 2.1.1). Theoretically, the FFI model can be extended to any number of choice-alternatives, if the inhibitory weights are adapted accordingly. This assumption might be plausible in an experiment where trials with different numbers of alternatives are present in separate blocks of trials. By contrast, if the different trials are randomly interleaved, as in the experiment we present in Chapter 5 and (Churchland et al., 2008), sufficient neural plasticity to adapt the connection weights is hardly practical in the short time between trials.

b) Lateral inhibition and the leaky competing accumulator

Apart from feedforward inhibition, also lateral inhibition between the two integrators, here the two LIP populations, could effectively render a diffusion process through competition within LIP (Fig. 2.6C). This picture

is consistent with the established assessment that long range cortical connections are mostly excitatory and inhibition thus acts locally within a cortical column (e.g. Lund et al., 2003). Usher and McClelland (2001) proposed a model called the "leaky competing accumulator" (LCA) model, which in its simplest form it can be written as:

$$dv_1 = (-kv_1 - wv_2 + I_1)dt + \sigma dW_1$$
$$dv_2 = (-kv_2 - wv_1 + I_2)dt + \sigma dW_2$$
$$, \quad v_1(0) = v_2(0) = 0.$$

(2.14)

Here, $k > 0$ is the decay rate, or leak, equivalent to $\lambda < 0$ in the Ohrnstein-Uhlenbeck model (Eq. 2.11), and w denotes the inhibitory connection strength between the integrator units. Usher and McClelland (2001) further incorporated nonlinearity in their model in the form of a threshold-linear activation function, which prevents that firing rates drop below zero. The LCA model accounts for correct and error RT distributions without the need of trial-to-trial, or within-trial variability.

Notably, the authors also addressed multi-choice decision-making with the LCA model and found that the model captures the positive log-linear relation between reaction time and number of alternatives, known as Hick's law (Hick, 1952).

To conclude, Bogacz et al. (2006) have demonstrated that for a particular parameter range, namely large and balanced leak and inhibition, the dynamics of the LCA and the FFI model approximate a one-dimensional linear diffusion process equivalent to the DDM with perfect integration. Moreover, the LCA and consequently also the basic DDM are approximated by a more physiologically plausible connectionist model with pooled inhibition and recurrent excitation, if self-excitation balances the decay and inhibition is strong (Fig. 2.6D). Yet, whether the brain actually works in a parameter regime of perfect integration has recently been called into question by a RDM study with time-varying evidence (Huk and Shadlen, 2005; Wong and Huk, 2008). For an accurate description of real neural dynamics, nonlinear attractor states that arise from strong recurrent connections might not be negligible.

2.2.3 Attractor models

Neural networks with interconnected neurons, such as the pooled inhibition model displayed in Fig. 2.6D, form nonlinear dynamical systems, whose long-term behavior is determined by "fixed points", or "steady states". These fixed points can be attractive or repellent and their existence depends on different parameters of the system, here for example the recurrent connection weights, or the inputs to the neural network. A useful analogy of the system's trajectory through state space is a particle

| Stable spontaneous state | Bistability | Multistability |
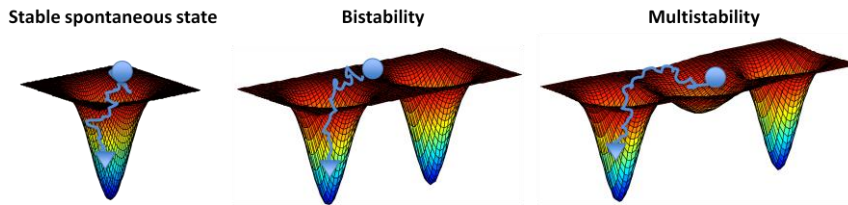
Fig. 2.7 Schematic of possible attractor configurations in the attractor network of binary decision-making.
Depending on the network parameters and inputs, one, two or three attractors can be simultaneously stable.

that moves on an energy landscape with valleys, corresponding to attractors or stable fixed points, and hills, corresponding to unstable, repellent fixed points (Fig. 2.7).

Decision-making can be implemented in an attractor network with at least two stable fixed points ("bistability"), representing the two choice-alternatives. The decision process is then regarded as the transition from an initial starting point towards one of the two attractors. Once such a decision-attractor is reached, this state will persist except for high levels of noise or perturbations and can thus be associated with persistent neural activity.

In the pooled inhibition model, decision-attractors emerge through strong recurrent connections, which form positive feedback loops of activity in the excitatory neural populations. Runaway activity is averted through negative feedback from the inhibitory neural population.

The biophysically realistic attractor model that forms the basis of our theoretical work, is a spiking-neuron implementation of the pooled inhibition model (Wang, 2002). Due to the nonlinear response properties of the spiking neurons, the full model can sustain a state of low spontaneous activity for a physiological range of background activity (Amit and Brunel, 1997). Therefore, depending on the amount of sensory inputs and the strength of the recurrent connections, the model can work in three different dynamical regimes: (1) only the spontaneous state of low firing is stable, (2) a bistable regime of categorical decision-making, and (3) a "multistable" regime, where the spontaneous and the decision attractors are stable (Fig. 2.7). In the multistable regime, transitions from the spontaneous state to a decision-attractor can happen due to noise fluctuations that are large enough to drive the system across the "hill", or unstable fixed point, into the "basin of attraction" of one of the decision-attractors.

a) Biophysically realistic attractor model with spiking neurons

Although the connectionist models discussed in Section 2.2.2 schematically describe neural processes in area MT and LIP, they still lack a direct connection between model variables and real neural parameters. This is different in the biophysically detailed implementation of the pooled inhibition model proposed by Wang (2002), where single neurons are modeled as leaky integrate-and-fire (LIF) neurons (Tuckwell, 1988) with conductance-based synaptic responses, described by realistic synaptic kinetics. The model was initially developed to account for (persistent) neural activity of PFC neurons during working memory tasks (Brunel and Wang, 2001). Its application to decision-making was inspired by the experimental observation that neurons, which exhibit ramping activity, characteristically show persistent neural firing in delayed memory or decision tasks (Gnadt and Andersen, 1988; Shadlen and Newsome, 2001). Wang (2002) successfully applied the spiking-neuron attractor model to behavioral and neurophysiological decision-making data measured from primates performing a binary RDM discrimination task (Roitman and Shadlen, 2002). In this context, the attractor network can again be viewed as the representation of a local microcircuit in area LIP (or PPC in general).

Physiological neural firing rates are obtained by averaging over the simulated action potentials, or output "spikes", of distinct neural populations of LIF neurons in the network. Each LIF unit is characterized by its subthreshold membrane potential

$$C_{\mathrm{m}} \frac{dV(t)}{dt} = -g_{\mathrm{m}}(V(t) - V_{\mathrm{L}}) - I_{\mathrm{syn}}(t),$$
(2.15)

with resting potential V_{L}, membrane capacitance C_m, membrane leak conductance g_m. I_{syn} is the total synaptic current flowing into the cell. When the membrane potential V of a LIF neuron reaches the firing threshold V_{th}, it is reset to V_{reset} and a spike is emitted to all connected neurons with a subsequent absolute refractory period of τ_{ref}. Accordingly, LIF neurons do not explicitly model action potentials, but give a realistic account of the subthreshold membrane potential. Excitatory synaptic currents between LIF neurons are mediated by fast AMPA and slow NMDA glutamate receptors, inhibitory synaptic currents by GABA$_{\mathrm{A}}$ receptors. The total synaptic current is given by the sum:

$$I_{\mathrm{syn}}(t) = I_{\mathrm{AMPA,rec}}(t) + I_{\mathrm{NMDA,rec}}(t) + I_{\mathrm{GABA}}(t) + I_{\mathrm{AMPA,ext}}(t).$$
(2.16)

The attractor network is organized into separate populations of LIF neurons, termed "pools", which share common inputs and connectivities (Fig. 2.8A). As in the connectionist version, the spiking-neuron model contains one homogenous pool of inhibitory neurons, globally connected

to all network neurons. The two integrator units are implemented by two "selective pools" of excitatory neurons, which are thought to respond *selectively* to one of the two possible directions of coherent motion and, hence, reflect the possible choice-alternatives in the RDM task (S1, S2 in Fig. 2.8A).
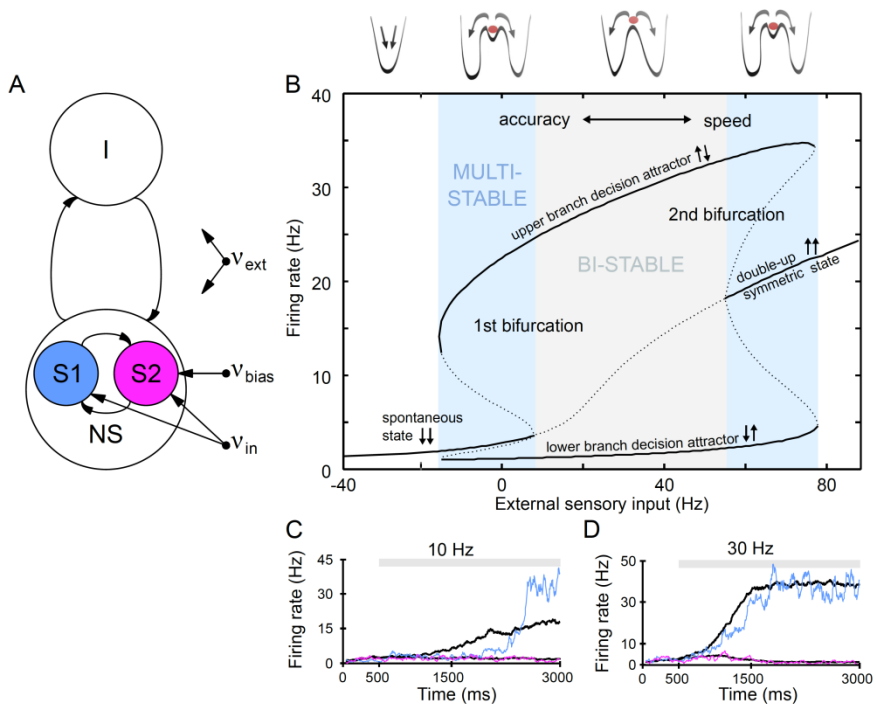
Fig. 2.8 Biophysically realistic attractor network of slow perceptual decision-making.
(A) Schematic representation of the network. The excitatory neurons are organized in three pools: the nonselective neurons (NS) and the two selective pools (S1, S2) that receive the inputs encoding each stimulus (with rate v_{in}). An additional bias (v_{bias}) can be applied to one of the two selective pools. All neurons also receive an input (v_{ext}) that simulates the spontaneous activity in the surrounding cerebral cortex. (B) Stable (solid lines) and unstable (dotted lines) fixed points depend on the external sensory input, They were calculated with the mean-field approximation of the network (Brunel and Wang, 2001). (C,D) Single trial (colored traces) and mean firing rate evolution (black, averaged over 20 trials) of the selective pools for different inputs, denoted by the gray bars ($v_{bias} = 0$). (C) Noise induced transition from spontaneous to decision state (low inputs, multistable regime). (D) Input-driven transition (bistable regime). Simulations were performed with a synaptic strength of $\omega_+=1.68$ within selective populations; all other parameters were taken from (Wang, 2002).

Moreover, a third excitatory pool of "nonselective" neurons represents activity of surrounding LIP neurons that are not selective to either direction.

All neurons in the network receive an external background input simulated by uncorrelated, stochastic Poisson spike trains applied independently to the individual neurons. This background activity determines the spontaneous firing rate of all network neurons and is the main source of noise in the network together with finite-size effects due to the limited number of neurons. Here, the term "finite-size noise" describes the variance in the recurrent inputs, which tend to zero when the size of the network is increased.

In addition to the background input, the selective pools further receive time-dependent external inputs, corresponding, for instance, to the sensory stimuli in the RDM experiment, the visual R-targets and the motion stimulus.

As explained above for the general case, decision-attractors emerge in the network due to the strong recurrent connectivity ω_+ of neurons *within* one selective pool, while the connections *between* the two selective pools are weaker than average $\omega_- < 1$[4]. Fig. 2.8B displays an example of a typical attractor landscape for strong ω_+ as a function of increasing external inputs applied equally to both selective pools.

Without any external sensory inputs to the selective pools (0 Hz), the system will naturally rest in its spontaneous state with similarly low firing rates in all excitatory neural pools.

If a sensory stimulus is applied to the model, which increases the external inputs to the selective pools sufficiently (> 10 Hz), the spontaneous state becomes unstable in a "subcritical pitchfork bifurcation" leading to bistability between the decision attractors (gray area in Fig. 2.8B). The network then operates in a region of categorical decision-making, where one selective pool will settle at the upper decision branch with high firing rate ("winner"), the other will decay to the lower branch ("loser"). In this case, the transition from the spontaneous state to the decision state is "input-driven" and can be gradual, in the order of several hundred milliseconds, even in single trials (Fig. 2.8D). These gradual transitions between attractor states, corresponding to the decision process, are a distinguishing feature of the biophysically realistic attractor model and rely on the slow kinetics of the NMDA receptors ($\tau_{NMDA,decay} = 100$ ms). Consequently, the network's behavior is not just dominated by its steady states, it also exhibits prolonged responses to momentary sensory inputs, with a characteristic time constant of up to a second,

---

[4] This is consistent with a Hebbian rule, as the activity of neurons that are selective for the same feature has supposedly been correlated in the past, while neurons encoding opposite directions rather fired in an anticorrelated manner.

during which the model effectively integrates the incoming inputs (Wang, 2008).

As depicted schematically in Fig. 2.7, also in the multistable regimes decision-making is possible. There, the transition from the spontaneous or symmetric state to a decision state is induced by noise fluctuations and can be rather sharp in a single trial (Fig. 2.8C). Yet, the trial-averaged activity builds up slowly, as observed experimentally in decision-related neurons (see 2.1.2). This ramping activity is, however, obtained in a conceptually different way compared to the bistable regime: here, the gradual build-up is an artifact of averaging across abrupt transitions at random times (Marti et al., 2008).

For sufficiently strong connection weights ω_+, as in the example of Fig. 2.8, the network can exhibit persistent activity, meaning that the high firing rates of the "winner" population can be sustained even after all external sensory inputs are switched off. This is because at 0 Hz, with only background activity, the system is already in the multistable regime, where the decision states are stable in addition to the spontaneous state. Under these conditions, the decision states would only destabilize if negative inputs were applied to the selective populations. Persistent activity is a characteristic feature of all biophysically realistic attractor models that we will explore in this thesis. Therefore, the bifurcation between the multistable regime at low inputs and the bistable regime is the first bifurcation with relevance to our purposes.

In the vicinity of this "first bifurcation", slow integration is enhanced above the intrinsically slow kinetics mediated by the NMDA receptors, as the effective time constant of the system exceedingly increases close to the bifurcation (Wong and Wang, 2006; Roxin and Ledberg, 2008). Therefore, in this dynamical region performance is high and reaction times are rather long, because of long stimulus-integration times. For this reason, previous analyses of the binary attractor model particularly concentrated on the dynamics in the proximity of this first bifurcation, where the spontaneous state destabilizes (Wang, 2002; Wong and Wang, 2006; Marti et al., 2008).

On this note, exploring the network dynamics at the other end of the bistable regime is one of the main objectives of this thesis. For sufficiently high external sensory inputs, the network again enters a multistable regime. Crossing the "second bifurcation", a symmetric "double-up" state becomes stable, where both selective pools fire with intermediate, elevated rates. This double-up symmetric state has recently been deployed to explain LIP responses to static visual stimuli, as for example the response targets in the RDM paradigm (Wong et al., 2007; Furman and Wang, 2008). Assuming high selective inputs with R-target onset, the high firing rates of LIP neurons prior to the motion stimulus (Roitman and Shadlen, 2002; Huk and Shadlen, 2005; Churchland et al., 2008; Kiani

and Shadlen, 2009) can be reproduced with the attractor model in the double-up state.

Across the bistable regime, in-between the two bifurcations, higher external inputs to both selective populations lead to faster reaction times and less accuracy, congruent with a speed–accuracy tradeoff (Wong and Wang, 2006; Roxin and Ledberg, 2008). This dependency of decision-behavior on the amount of unbiased external sensory inputs to both selective populations explicitly arises from the nonlinearity of the attractor model. In linear models, such as the DDM, changes in the common sensory evidence to both decision alternatives would not affect decision behavior. This is obvious in the case of the DDM, as the DDM characteristically accumulates only the evidence difference of the two alternatives. We will see in Chapter 3 that input-dependent decision behavior might provide means to distinguish between linear and nonlinear modeling approaches.

As a final note, the decision-attractor would automatically provide an upper bound for the neural activity of the winning selective population. Nevertheless, the model's decision is typically determined by a fixed firing rate threshold independent of the applied amount of sensory inputs, in line with neurophysiological evidence from LIP neurons (see 2.1.2). How this decision threshold is read out or adjusted by down-stream areas is not explicitly included in the attractor model. Yet, possible extensions have been suggested, which implement the decision threshold involving cortico-collicular and cortico-basal ganglia circuits (Lo and Wang, 2006; Bogacz and Gurney, 2007).

Taken together, the characteristic features of the biophysically realistic spiking-neuron attractor model are:

- strong recurrent connections within the selective neural populations, which generate attractor states,

- global feedback inhibition enabling winner-take-all competition,

- stochasticity because of finite-size effects and random Poisson inputs to the network,

- a long synaptic time-constant (NMDA) facilitating the integration of incoming neural activity.

A detailed mathematical description of the neural dynamics is given in the Appendix (A.1), together with all parameters and specifications of the different model variants and extensions that we implemented in this dissertation.

b) Model reductions

Simulating populations of individual and realistic neurons as described above is necessary to simulate realistic neuronal dynamics,

physiological responses and behavior. Nevertheless, to gain an analytical understanding of the population dynamics, a reduced, mathematically tractable description can yield deeper insights into the essential, collective model behavior.

Mean-field approximation. Taking a mean-field approach, Brunel and Wang (2001) considerably reduced the state variables of the network by replacing the different populations of spiking neurons with an approximation of their average population activity. Because of the recurrent connections in the network, the population firing rates have to be determined self-consistently based on the input currents to the neural pools, which in turn depend on the firing rates. Equalizing the pre- and postsynaptic activity, the possible fixed points or steady states of the population firing rates can be obtained.

Several approximations have to be assumed in order to arrive at a closed system of one nonlinear equation for each population in the network. First, postulating that individual neurons fire spikes independently, according to a stationary Poisson process, the net input is treated as a Gaussian random process. This assumption generally holds in the limit of infinitely large networks, where each neuron receives a large number of presynaptic currents, which each deflect the membrane potential only minimally compared to the voltage distance between the resting and threshold potential. Second, only fast AMPA-mediated external inputs are assumed to contribute to the fluctuations in the synaptic current. Fluctuations in the NMDA and GABA currents are neglected as they are supposed to filter out, due to the longer synaptic time constants.

Finding a self-consistent solution for the population rates is further complicated by the nonlinear properties of the NMDA receptors. Therefore, NMDA saturation is approximated by calculating the average NMDA gating variable as a nonlinear function of the presynaptic rate. In addition, the voltage dependence of the NMDA conductance is linearized around the average neural potential.

The final set of mean-field equations obtained through the above approximations is given in Appendix A.1.2. For a detailed mathematical derivation we refer to the original publications (Brunel and Wang, 2001) and (Renart et al., 2003). Solving the mean-field equations is computationally much less intense than running simulations with the full spiking network. The mean-field analysis thus allows calculating the steady state firing rates of the attractor model for a wide range of parameters. This makes it feasible to scan the parameter space in order to find a parameter set matching experimental findings.

In sum, by solving the mean-field equations for a set of initial conditions (here the initial firing rates of each neural population) one obtains the approximated average firing rate of each pool when the system

has settled into a stationary state after the period of dynamical transients. The mean-field reduction, however, does not provide an accurate description of the temporal dynamics.

Two-dimensional reduction. Based on the mean-field approach, Wong and Wang (2006) further reduced the biophysically realistic spiking-neuron model to a two-variable system. In particular, they fitted a simplified input-output function to the complex first-passage time formula used in the mean-field to describe the output firing rates as a function of the mean synaptic inputs. With the assumption that the network dynamics are dominated by the slow NMDA receptors, they further set the firing rate of the nonselective pool to a constant mean firing rate (2 Hz) and linearized the input-output relation of the inhibitory neurons. Thereby, inhibition could be incorporated into the selective populations as mutual negative inputs. The nonselective and interneurons could thus be eliminated, leaving two neural units with self-excitation and effective mutual inhibition.

The two-dimensional reduction is particularly useful to perform phase-plane analyses in order to elucidate the different dynamical regimes of the network. To simulate the noisy temporal evolution of the spiking neural network, Wong and Wang (2006) explicitly added a noise term to the external inputs, described by an Ohrnstein-Uhlenbeck process (white noise filtered by a short, AMPA synaptic time constant). The two-dimensional reduction can thus be viewed as a closely related connectionist version of the full spiking model. In this way, it can also account for decision-related behavior and neural activity, albeit without explicit analogy to real neural parameters (Wong et al., 2007).

Nonlinear diffusion. Instead of a two-component system of rate equations as in (Wong and Wang, 2006), Roxin and Ledberg (2008) derived a one-dimensional *nonlinear* diffusion equation to describe the asymptotic behavior of winner-take-all models in the proximity of the bifurcation to bistability, where the spontaneous state destabilizes. Their reduction is universally valid for all winner-take-all models, but also allows to relate the variables of the nonlinear diffusion process to those of the full spiking-neuron model and thus to neurobiologically meaningful quantities. In the last section, we already mentioned a particularly relevant prediction based on this nonlinear one-dimensional reduction, namely, that the speed-accuracy tradeoff can be implemented by changes in the common inputs to both selective neural populations, instead, or in addition, to an adaptation of the decision threshold. We will confirm this relation between selective inputs and the speed-accuracy tradeoff with the full spiking model in Chapter 3 (see in particular 3.4.2).

2.3 Distinguishing model approaches[5]

As we have seen, models on the accumulation of noisy evidence, as for instance in the random-dot motion paradigm, come in a huge variety of flavors. Although they differ in fundamental features, such as network structure and connectivity, in practice, it may be very difficult to distinguish between them on the basis of just behavioral data or mean firing rates (Bogacz et al., 2006; Ditterich, 2010).

Ratcliff and Smith (2004) evaluated four types of sequential sampling models and the LCA model against three sets of psychophysical 2AFC experiments. In particular they compared the three models we presented in Section 2.2.1, the DDM, the O-U and the race model, and a so-called "Poisson counter"[6] model (Townsend and Ashby, 1983), all with trial-to-trial variability in drift, starting point and non-decision time.

Of all models considered, only the Poisson counter model failed to match the empirical data and faster mean errors still resulted problematic for the race model. The Poisson counter model also proved inferior to the DDM when compared to the neural activity of superior colliculus build-up neurons from macaque monkeys performing a 2AFC task (Ratcliff et al., 2003). The activity pattern predicted by the DDM, however, resembled the observed neural firing rates, suggesting that build-up cells in the superior colliculus might participate in a diffusion-like decision process.

Because of the mutual mimicry between models (Ratcliff and Smith, 2004; Bogacz et al., 2006), finding new analytical methods and intelligently designed experiments to distinguish the different approaches is a major future challenge in the field of perceptual decision-making.

One approach along that line was conducted by Huk and Shadlen (2005). By adding brief motion pulses to a standard RDM stimulus, they first of all provided strong physiological support for a temporal integration in LIP. However, their findings revealed a departure from *perfect* integration, as the effect of the motion pulse decreased with its onset time. Later motion pulses thus influenced behavior and neural activity less than earlier motion pulses. Neither a perfect DDM, nor leaky integrators could reproduce this experimental finding, while the time-varying dynamics of the attractor model explained both behavioral and neural data (Wong et al., 2007; Wong and Huk, 2008). Still, time-varying effects such as

---

[5] Part of the review presented in this section is adapted from a discussion published in Masquelier T, Albantakis L, Deco G (2011) The timing of vision - how neural processing links to different temporal dynamics. Front Psychol 2:151.
[6] The Poisson counter model resembles the race model, with the difference that evidence is counted in discrete units, delivered at random times, with exponentially distributed intervals. Therefore, it can be interpreted as an independent accumulation of two spike trains.

decreasing decision bounds or an "urgency" signal might produce decreased sensitivity to later perturbations also in the DDM and LCA.

Recently, also multiple-choice decision task received increasing attention in the context of model distinction (Leite and Ratcliff, 2009; Ditterich, 2010; Purcell et al., 2010; Churchland et al., 2011).

Analyzing higher-order statistical properties of neurophysiological data from their 2- and 4-alternative RDM task, Churchland et al. (2008; 2011) were able to distinguish between models categorized by their different sources of variability. Models with just one source of variability, such as the LATER model (2.2.1b) and a model by Cisek et al. (2009) with fixed slope, but a random distribution of firing rates at each time-step, failed to account for the higher-order measures, although they agreed with behavior and mean firing rates. On the other hand, all different model implementations of a stochastic accumulation-to-threshold tested in Churchland et al.'s (2011) study could account for variance and within-trial correlations, in addition to behavioral data and first-order firing rates. In particular, the tested models included the drift–diffusion model (Ratcliff and Rouder, 1998), a model based on probabilistic population codes (Beck et al., 2008), and the reduced version of the attractor model by Wong et al. (2007).

Based on human behavioral data from a RDM task with three alternatives and three motion components, Ditterich (2010) intended to distinguish more detailed aspects of conceptual accumulation-to-bound models with regard to their goodness of fit and their neurophysiological predictions. Perfect integrators were compared to leaky, saturating integrators, with either feedback or feedforward inhibition. As we have seen, in the case of two alternatives, most of the discussed models proved equivalent to the DDM for certain parameter ranges (Bogacz et al., 2006). Therefore, it might not be too surprising that none of the models could be excluded based only on the fits to behavioral data of a 3-alternative RDM task (Niwa and Ditterich, 2008). Yet, the models differ substantially in their neurophysiological predictions on how the integrator states should evolve over time (see Table 2 in Ditterich, 2010). Invasive neural recordings from monkeys performing the same task will hopefully soon settle the dispute. Moreover, feedforward and feedback inhibition respectively suggest either negative or positive correlation between the integrator units, which might be tested with multi-electrode recordings. Finally, in the case of an equal amount of motion coherence in all three directions, Niwa and Ditterich (2008) measured faster mean reaction times for higher coherence levels. While models with feedforward inhibition require a scaling of the variance of the sensory signals in order to account for this effect, conceptual models with feedback inhibition could explain the result just with a change of the mean input (Ditterich, 2010).

Considering all the current evidence presented in this chapter, so far two types of decision-making models have proven particularly successful: on the one hand, the extended drift-diffusion model and its connectionist implementations account for a vast range of behavioral data. They also conceptually represent neural activity during the decision-making period. On the other hand, Wang's (2002) physiologically-detailed attractor model and its reductions (Wong and Wang, 2006), which mimic real neural dynamics, accurately simulate behavioral data and LIP activity during the RDM task. Moreover, they account for persistent activity and the nonlinear, time-dependent effects of motion pulses.

In the following three chapters, we intend to confirm the accordance of attractor dynamics to higher-level cortical processes in general, and changes of mind and multiple-choice decision-making in particular. What is more, we will compare extensions of the attractor model to other, more phenomenological, modeling approaches, and propose new ways to distinguish between them.

3 CHANGES OF MIND IN AN ATTRACTOR NETWORK OF DECISION-MAKING

The work presented in this chapter is published in PLoS Computational Biology.[7]

3.1 Introduction

Some of the decisions we have to make in our lives are irretrievable. Once an alternative was chosen, it cannot be undone. Often, however, our decisions are not strictly binding. We might be granted a second chance to decide and a previous choice can then be adjusted if we have changed our minds.

As we have seen in the last chapter, the decision process is traditionally regarded as a decision variable evolving in time, until a termination criterion is reached. Correspondingly, the lateral intraparietal (LIP) cortex was identified as a possible candidate for a neural decision variable. This is due to the fact that firing rates of LIP neurons gradually increase during motion-viewing in the RDM task and correlate with subjects' choices and reaction times (2.1.2). Moreover, recordings from LIP neurons provide evidence for a fixed decision threshold: if the monkey chose the target in the response field of the recorded neuron, the firing rate variation across trials reached a minimum shortly before saccade onset (Roitman and Shadlen, 2002; Churchland et al., 2008). This means that on each trial about 80 ms before the saccade, the neurons fired at an approximately stereotyped rate, indicating a decision threshold. The chosen motor response, typically a saccade, then marked the end of the decision-trial.

With this established concept of a decision threshold, how could a change of mind be induced after a first decision was already made? Or more generally: What happens in our brains if we change our mind?

To elucidate these questions, Resulaj et al. (2009) developed a psychophysical RDM task, where human participants had to indicate their choice by moving a handle towards a left or right target (Fig. 3.1A). Because this hand movement is continuous, contrary to ballistic saccades or pressing a button (Palmer et al., 2005), occasionally changes of mind could be observed directly by recording the handle traces. Changing improved the overall accuracy, but depended on task difficulty: most

---

correcting changes, meaning changes from the incorrect to the correct side, were observed at intermediate levels. Erroneous changes from the correct to the wrong side, however, increased monotonically with difficulty.

These findings pose a challenge for attractor models. This class of models implements decision-making by diffusion in a nonlinear landscape of stable fixed points, which act as decision-attractors. Once a decision-attractor is reached, this state will persist except for high levels of noise or perturbations and is thus rather counterintuitive to a change of mind. On the other hand, due to the stable attractors, those models account for persistent activity frequently observed in decision-related neurons (see 2.2.3).

In this chapter we show that changes of mind (after a first decision) are entirely consistent with attractor dynamics. In particular, they arise naturally during the itinerant transients following sensory perturbation, if the system lies close to a bifurcation (or phase boundary) that separates a neuronal state of categorical decision-making from a multi-stable region. There, the decision process is impeded by a second attractor, where both populations encoding the possible alternatives fire at high rates. This facilitates changes of mind. Moreover, by replicating the psychophysical data of Resulaj et al. (2009) with a biophysically realistic attractor model with spiking neurons, we gained neurophysiological predictions on neural firing rates during the change process. In all, our results offer testable predictions on the attractor concept and general principles of decision-making like the speed-accuracy tradeoff and a fixed decision threshold.

Our main results are the following:

- Despite their fixed-point stability, attractor models can account for changes of mind.

- Low decision thresholds and high incoming activity (speed pressure) favor changes.

- The model fits the experimental data best close to a bifurcation point at high inputs.

- We predict a switch in neural activity during changes of mind, which might be indicative of the neural decision threshold.

- We further suggest that the brain operates over the whole range of inputs, which enable decision-making.

Fig. 3.1 Experimental design, network architecture and stimulation protocol.
(A) RDM paradigm with manual indication of choice as in Resulaj et al. (2009).
(B) Diagram of the binary attractor model for decision-making (Wang, 2002).
Unlabeled arrows denote a synaptic weight of 1. The model is fully connected.
(C) Time course of target and motion input to the selective populations in order
to model the experimental design of the RDM task.

3.2 Methods

With the objective to gain understanding of the actual brain processes during changes of mind, in the following we apply a biologically-inspired attractor model of decision-making to the psychophysical findings of Resulaj et al. (2009). The general network structure of the biophysically realistic attractor model we present here (Fig. 3.1B), is identical to the two-alternative attractor model first introduced by X.J. Wang (2002) and described in Section 2.2.3. The specific network parameters and inputs, however, were adapted in order to account for changes of mind. Table A.1 and A.2 with all default parameter values and the full network details can be found in the Appendix A.1.3, together with specification of the mean-field analysis. The simulation parameters of the diffusion model we used for the model comparison in Section 3.3.6 are described in Appendix A.1.4.

3.2.1 Experimental paradigm

The experimental task sequence used by Resulaj et al. (2009) is illustrated in Fig. 3.1A. Three human participants were tested. While the participants were holding a handle at the starting position, a patch of randomly moving dots appeared after a random delay (0.7-1.0 s). Depending on the trial difficulty, a certain percentage of these dots were moving coherently to the left or right. The subjects had to decide within 2 s on the net direction of dot-motion and to report their choice by moving the handle in the corresponding direction towards a response target (R-target, Fig. 3.1A, red dots). They were asked to respond as quickly and accurately as they could. Once they initiated the hand movement, they had to reach the R-target within a time limit of 700 ms.

In the majority of trials the subjects moved the handle directly to one of the R-targets. Some trajectories, however, revealed a change of mind during the movement: they started towards one direction but terminated at the opposite R-target. Importantly, the moving-dot display was switched off when the handle left the starting position. Participants thus occasionally changed their mind on the way towards the R-target, although the motion stimulus was no longer visible.

Three dependent variables were evaluated in the psychophysical experiment: (1) reaction times (RTs), which correspond to the time of movement initiation, when the handle left the starting position; (2) the participants' accuracy, or "performance"; (3) the probability of a change of mind, obtained from the hand-movement trajectories.

For trials including changes of mind, the *initial* performance is given by the initial direction of the hand movement, while the *final* performance is determined by the finally selected R-target. Moreover Resulaj et al.

(2009) distinguished between *correcting* changes of mind, initially directed towards the incorrect R-target, but then changed to the correct R-target, and *erroneous* changes, which turned from the correct to the wrong choice.

3.2.2 Attractor network for changes of mind

a) Network structure

In brief, the network consists of two subpopulations (pools) of excitatory pyramidal neurons, which implement the two decision alternatives (Fig. 3.1B, red). Each of them is selective for one of the two target directions. A nonselective excitatory population represents activity of surrounding neurons that are not selective to either direction. Competition arises in the network due to global feedback inhibition by a population of inhibitory neurons. To accurately simulate neural activity, the 1,000 network neurons are modeled as integrate-and-fire neurons with synaptic currents mediated by AMPA, NMDA and $GABA_A$ receptors with biophysically realistic conductances and time constants (Table A.1). The strong recurrent connections between neurons from the same selective pool (ω_+), together with the long synaptic time constant of the NMDA currents enable the model to integrate incoming activity and sustain elevated firing rates for several hundred ms.

b) Network inputs

During the simulation, each neuron individually receives stochastic excitatory Poisson inputs from several external sources. The noise fluctuations around the mean external input applied to each neural population thus depend on the amount of neurons in the respective pool and would be zero for an infinite number of neurons ("finite size" effect). For the two selective populations (consisting of 160 neurons in the present network) the standard deviation is 17 Hz given a total external input of about 2.4 kHz (see below 3.3.3). These 2.4 kHz, equal to 800 afferent neurons firing at 3 Hz, simulate the spontaneous activity in the cerebral cortex outside the local network.

On top of this background activity, an external target and motion input are applied to the selective neural populations only (Fig. 3.1C). They correspond to the sensory stimuli during the RDM experiment: the visually shown R-targets and the random-dot motion respectively.

Target input. In Resulaj et al.'s (2009) experiment the two possible R-targets were visible throughout the trial. During neurophysiological single cell recordings combined with the RDM task, one target is always placed in the response field of the recorded LIP neuron. Thus, the selective populations are supposed to respond not only to the motion evidence in favor of one of the two target directions, but also to the R-

targets themselves. The time course of the target input aims to replicate the evolution of LIP firing rates after R-target presentation. In previous neurophysiological studies, LIP firing rates have been found to rise steeply with the appearance of the possible targets, followed by a "dip" in activity at the onset of the motion stimulus (Huk and Shadlen, 2005; Churchland et al., 2008; Kiani and Shadlen, 2009). Correspondingly, in the simulations the target input (Fig. 3.1C, red) is composed of initially high activity with a subsequent decline of input activity. Specifically,

$$
\nu_{\text{target}} = \begin{cases} 0 \text{ Hz} & 0 < t < t_{\text{target}} + 100 \text{ ms} \\ (350 + 100\,exp(-(t - t_{\text{target}} - 100 \text{ ms})/\tau_1)) \text{ Hz} & t_{\text{target}} + 100 \text{ ms} \leq t < t_{\text{motion}} + 80 \text{ ms} \\ (85 + 265\,exp(-(t - t_{\text{motion}} - 80 \text{ ms})/\tau_2)) \text{ Hz} & t \geq t_{\text{motion}} + 80 \text{ ms} \end{cases}
$$

(3.1)

with t_{target} = 400 ms and t_{motion} = 1,300 ms plus an assumed latency of 100 ms and 200 ms, respectively, for the signals to arrive in area LIP (Churchland et al., 2008). The initial exponential decay τ_1 = 100 ms can be explained by short term adaptation. Due to the exponential decrease of the target input with τ_2 = 15 ms, starting with a latency of 80 ms after motion-stimulus onset, the target input is already decaying for 120 ms, before the motion input arrives in LIP with a latency of 200 ms. This causes the dip of firing rate in the simulations. Physiologically, this dip might be explained by an attentional shift or upstream inhibition of the R-target signal with the onset of the motion-stimulus (Wong et al., 2007; Furman and Wang, 2008).

Note that the specific parameters of the target input are irrelevant as long as, first, the initial inputs are high enough to shift the network from the spontaneous to the symmetric state with high firing rates in both selective populations and, second, the target input is reduced sufficiently with motion onset to allow competition (see attractor landscape in Fig. 2.8 and Fig. 3.6). What is more, the model is generally capable of decision-making and changes of mind even in the absence of a target signal (Fig. 3.A.1).

Motion input. The motion input represents activity of middle temporal (MT) area neurons projecting to PPC. MT neurons fire dependent on the amount of coherent motion towards their preferred direction (Britten et al., 1993). Accordingly, the different motion coherence levels are translated into a positive bias of motion input to one of the selective populations, balanced by a motion-input reduction in the other:

$$
\nu_{1,2} = \nu_{\text{motion}} \left(1 \pm \frac{c}{100\%}\right),
$$

(3.2)

with a time invariant rate of ν_{motion} = 70 Hz for 0% coherence. Thus, for 0% coherence in the RDM stimulus, both selective pools receive the same

amount of motion input (70 Hz, Fig. 3.1C blue), while for 100% coherence only one pool would receive the maximum motion input (140 Hz). We simulated six coherence levels: $c = 0\%$, 3.2%, 6.4%, 12.8%, 25.6%, and 51.2%.

In the following we refer to both the target and motion input as "selective inputs". Both selective inputs are present until the end of the simulation (3,500 ms).

c) Decision threshold and simulated changes of mind

With the start of the motion input the system dynamically evolves towards the decision state, where one of the two selective pools fires at a high rate, the other at a low rate. During this transition, a (first) decision is made when one of the firing rate transients crosses the decision threshold (44 Hz) with the additional condition that the difference between populations is at least 10 Hz. A trial was considered a change of mind, if the firing rate of the initially losing selective pool exceeded the (same) decision threshold after the first pool crossed, and their rates differed again by 10 Hz or more. Our main motivation to use a difference criterion in addition to the fixed threshold was to avoid very occasional joint threshold crossings to count as decisions (see example in Fig. 3.A.2E). As fluctuations in the firing rate of the selective populations are rather anticorrelated because of the global feedback inhibition and typically larger than 10 Hz, given the amount of noise present in the network, that constraint has only little effect on the simulation results.

d) A time-out for changes of mind

The motion stimulus in the experiments was turned off when the handle left the starting position. During the time of motor preparation and initiation, new evidence could already have arrived in LIP that was not taken into account for the first decision (~ 180 ms, Snyder et al., 1997; Cui and Andersen, 2007). In addition, the last evidence shown to the subject would reach LIP only after a sensory latency of about 200 ms (Roitman and Shadlen, 2002; Churchland et al., 2008). In total, after the first decision, new, yet unprocessed evidence on the motion direction, was possibly available to LIP for a time equivalent to the non-decision time $t_{ND} = 380$ ms of a trial, i.e. for the duration of motor initiation, plus the latency for the evidence to arrive in LIP. The assumed t_{ND} value of 380 ms for the non-decision time is in agreement with the fit of a simple accumulation-to-bound model to the experimental data of the three participants (Resulaj et al., 2009). Resulaj et al. (2009) indeed found, that random fluctuations in the motion stimulus during this time period correlated with changes of mind, indicating that the new evidence caused the subjects to change. In the model, a change of mind without motion input is very unlikely (see 3.3.5). As we were interested in the further

progression of the transients to the attractor states and for computational and analytical reasons, the motion input in the model lasted until the end of the trial simulation (3,500 ms). Therefore, we imposed a time-out of t_{ND} for changing after the first threshold crossing, which implements the experimental time limit for new evidence, caused by switching off the motion stimulus at movement initiation. Note that the simulations are still perfectly congruent with the experiment up to the first threshold crossing plus $t_{ND,}$ and also thereafter, as neither in the model nor in the experiment further changes (or threshold crossings) are expected.

The robustness of the model simulation to variations in the decision criteria and the non-decision time is shown in Fig. 3.A.2.

3.3 Results

3.3.1 Comparison to behavioral data

Generally, the simulated reaction times and percentages of correct choices fit the experimental results well for the applied amount of sensory inputs. This can be appreciated in Fig. 3.2 where the model's simulated behavior is compared to one participant from Resulaj et al. (2009)[8]. In the experiment the reaction time was set by the initiation of the hand movement. Accordingly, the simulated reaction time is composed of the time of first threshold crossing, plus the non-decision time $t_{ND} = 380$ ms.

Moreover, the model also replicated the frequency of changes observed experimentally (Fig. 3.2, right panel). Taking the changes of mind into account improves the performance, as correcting changes from wrong to correct choice are more frequent for all coherence levels, but especially for intermediate motion strengths (Fig. 3.2A, left panel, red line). Erroneous changes to the wrong alternative, however, decayed monotonically with increasing motion coherence. They are most frequent for low motion strengths and do not occur for high motion coherence.

In comparison to the experimental findings, the model predicts slightly more changes to correct and less to the wrong choice, which also explains the larger difference of performance with and without changes (see 3.4 Discussion).

Resulaj et al. (2009) further noted that a seemingly optimal strategy to opt for or against a change would be to always wait until the end of t_{ND} after the first decision and, thus, to consider all possibly available evidence. This, however, was not consistent with their experimental observations.

---

[8] See (Resulaj et al., 2009) for further comparison with their other two participants.

Along that line, we analyzed the time distribution of changes of mind in the attractor model (Fig. 3.3). In the simulations, the changes are broadly distributed across t_{ND}, with the exception that hardly any changes occur during the first 50 ms after the first decision. The distribution peak depends on the motion coherence level, with earlier changes for higher coherences (Fig. 3.3B). Interestingly, the time difference between threshold crossings for erroneous changes is not considerably shorter than for correcting changes, although there is more evidence in favor of changing in the case of an initially wrong choice. Erroneous changes just become overall less frequent with increasing coherence.

Moreover, in the simulation in at most 1.6% of the trials two changes occurred during t_{ND} (Fig. 3.2 right panel, dashed line). The second change was then neglected. Notably, these double-changes were indeed occasionally found in the experiments (M.N. Shadlen, personal communication). In summary, although we did not aim for a perfect quantitative fit to the experimental data, the psychometric functions
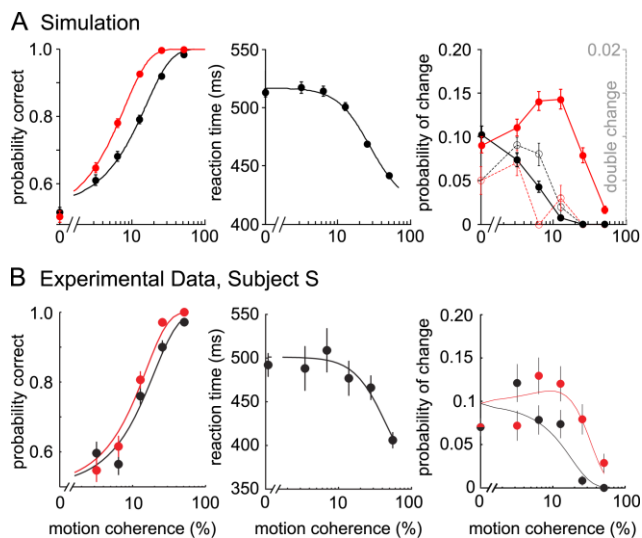
Fig. 3.2 Simulated psychometric functions, RTs and rates of changes compared to experimental data.
(A) Simulation data. For comparison, the experimental performance of Subject S from Resulaj et al. (2009) is shown in (B). (Left panel) Initial performance (black trace corresponding to choice at movement initiation) and final performance (red trace, corresponding to the finally chosen target). (Middle panel) Reaction times. (Right panel) Changes of mind. Erroneous changes are displayed in black, correcting changes in red (solid lines). Double changes in the simulations are shown on a ten times smaller timescale (right) (open circles, dashed lines). Black (red): proportion of erroneous (correcting) changes that switched a second time. Simulated psychometric functions were fitted by a logistic function (Eq. A.34), RTs by a hyperbolic tangent function (Eq. A.35).

obtained by our model simulations match the experimental observations very well in all relevant aspects.

3.3.2 Predictions on neural activity

Due to its biophysically realistic properties, the theoretical model offers predictions on neural firing rates during changes of mind. In Fig. 3.4A and E single trial examples of network simulations are displayed with and without changes of mind. In the trials with identical inputs to both selective pools (0% motion coherence), the decision which population activity will rise or decay is stochastic due to the Poisson inputs and finite-size noise fluctuations. The general temporal structure of the network activity matches single neuron recordings of primate LIP
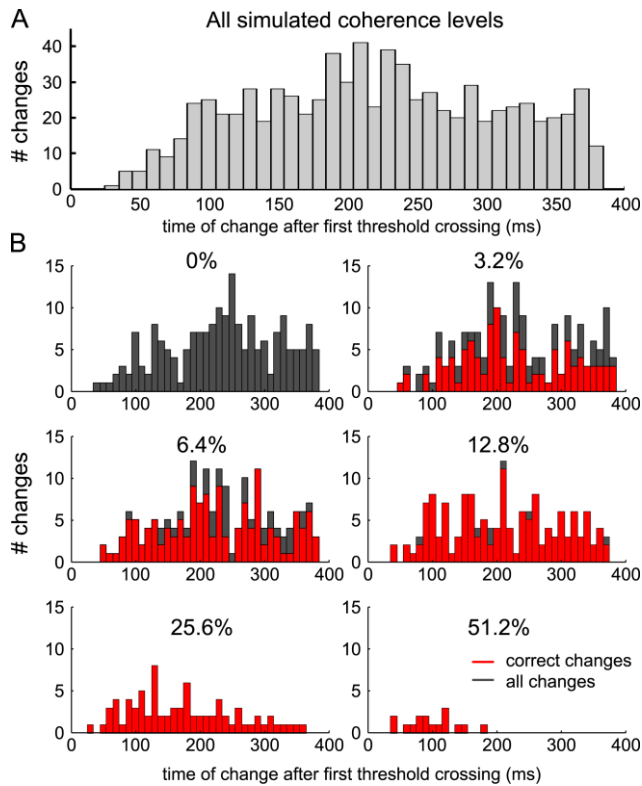
Fig. 3.3 Distribution of change times.
(A) Histogram of the time difference between the first and second threshold crossing (change of mind) for all change trials. The change times are broadly distributed from about 50 ms after the first decision to the time-out t_{ND} for changing. (B) Same as (A) separated into coherence levels. All changes are shown in dark grey. The correcting changes are overlaid in red, except for 0% coherence, where changes are neither correcting nor erroneous.

neurons with a high response to the target signals (from 500 to 1,300 ms), a subsequent dip of activity and a build-up of the firing rate after the onset of the moving dots (1,500 ms) (Huk and Shadlen, 2005; Churchland et al., 2008; Kiani and Shadlen, 2009). After an initial joint build-up during which the transients compete for the higher attractor state, the slope of the ramping activity is steeper with higher motion coherence (Fig. 3.4C,D average of correct trials at first threshold crossing).

Except for the highest motion coherence, this firing rate build-up is biphasic: after an initial steep increase independent of motion strength, the slope of the ramping activity decreases with lower motion coherence. To obtain sufficient changes of mind in the model simulations, the decision threshold was set relatively close to the divergence of the mean build-up activities for different motion coherences, which led to rather small differences in reaction times between the easiest and more difficult trials (see 3.4 Discussion). Nevertheless, the firing rate slopes clearly diverge with motion strength already before the threshold is reached (Fig. 3.4D).

In Fig. 3.4F we averaged all simulation trials with changes of mind, aligned to the first threshold crossing, which, if a constant non-decision time is assumed, corresponds to aligning to reaction time in the experiments. Thus, we show that the predicted rise and fall of activity during changes of mind might actually be observed experimentally, even if neural activities obtained in single cell recordings need to be averaged over trials to obtain reliable firing rates. In fact, even for a normally distributed non-decision time with moderate standard deviation, the switch in firing rates should still be discernible in neurophysiological experiments (see Fig. 3.A.2H).

3.3.3 Input fluctuation analysis

In the last section, the model's output firing rates were aligned to the first threshold crossing in order to predict neural activity during changes of mind. In the same way, the average firing rate input to the network after the first threshold crossing can give insights about the size of fluctuation necessary to evoke a change of mind.

In the simulations as well as in the experiment the mean coherence level, and thus the net evidence for or against one direction, does not change during a particular trial. Nevertheless, as most of the dots in the experimental RDM stimulus are moving randomly, the actual momentary level of coherent motion towards one target direction fluctuates around the set mean coherence. A measure of these stimulus fluctuations with respect to the monkeys' choices, the "motion energy", was found to support the initial decisions as well as the change of mind (Resulaj et al., 2009). More precisely, the fluctuations in the first 150 ms after stimulus onset acted as additional evidence in favor of the first decision (positive motion energy).
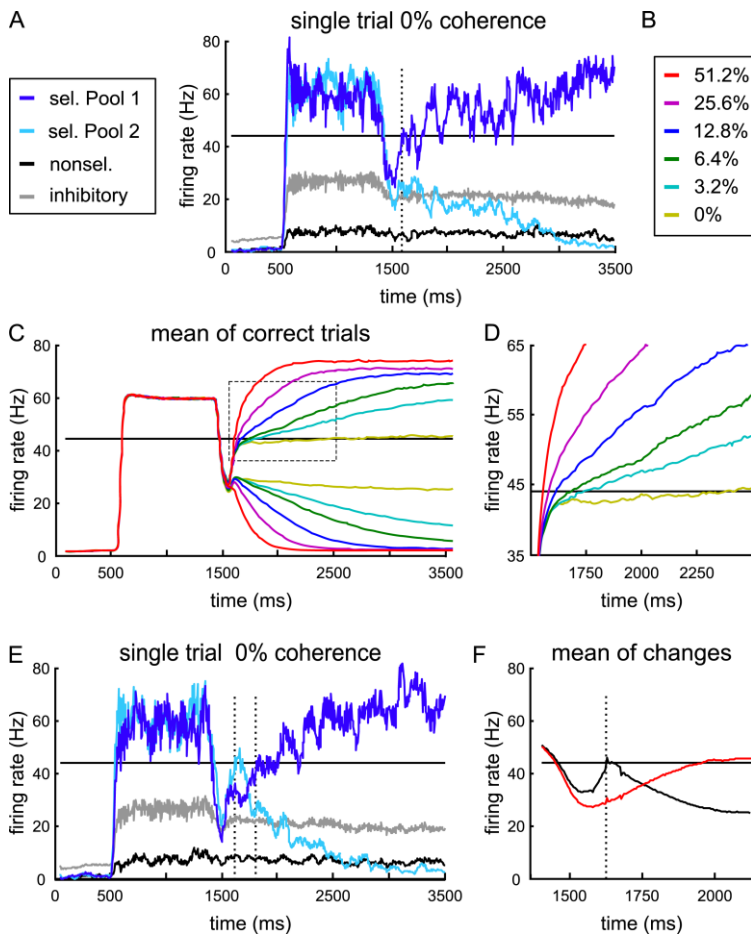
Fig. 3.4 Model prediction of LIP firing rate.
(A,E) Simulated temporal evolution of population-averaged firing rates for single trials. The dotted lines mark times of threshold crossings. The black line at 44 Hz indicates the threshold. (A) Example for a regular trial without change. (C, D) Mean of correct trials from 1,000 network simulations, shown for all motion coherences (Color code according to B). For each motion strength the firing rates were averaged according to the "winners" and "losers" of the first decision. (D) Blow up of dotted rectangle from (C). (E) In some cases the initially winning population (first threshold crossing) is overtaken by the other transient, which is counted as a "change of mind" trial. (F) Mean of all trials with changes (correct and error trials, all motion coherences) aligned to the first threshold crossing (dotted vertical line). Black: initially winning selective pool, red: finally winning selective pool.

In change trials the motion energy subsequently became negative, indicating that stimulus fluctuations played a causal role in switching through weakening or even reversing the preceding evidence, which favored the initial choice.

In the model simulations, the Poisson noise around the mean input rate corresponds to the experimental stimulus fluctuations. Fig. 3.5 shows the variation from mean input difference of the selective populations aligned to first threshold crossing and changes of mind (insets). In line with the experimental motion energy, the average input fluctuations across all change trials became negative after the first threshold crossing. Input fluctuations thus act as evidence against the initial choice. Note however, that for high coherence levels the changes do not depend on random fluctuations of the input, since it is mostly initial errors that are reversed by the designated input bias to the correct selective population. Interestingly, the fluctuation strength necessary to reverse a decision is in general not substantially higher than that causing the initial decision.

3.3.4 Mean-field analysis indicates proximity to bifurcation

As seen above, input fluctuations contribute to changes of mind in the attractor model. Whether such a fluctuation in incoming activity is large enough to elicit a change of mind, however, depends on the working point of the network in the attractor landscape. Keeping all other network parameters fixed, the working point can be adjusted by the mean inputs to the selective population in the network. While the model can match the experimentally obtained reaction times and performances for a large range of selective inputs, if the threshold is adapted accordingly (Fig. 3.A.3), the feasible range of network inputs is greatly reduced by the additional constraint to match the changes of mind.

Using a mean-field approximation of the model (Brunel and Wang, 2001), we analyzed the dynamical behavior of the network as a function of the selective input amplitude for the parameters that fit the changes of mind (see 2.2.3 and A.1.3). Simulating populations of individual and realistic neurons as described above is necessary to simulate realistic neuronal dynamics, physiological responses and behavior. However, to understand the underlying attractor and dynamical structures prescribing the behavior of population dynamics, we had to use a simpler model that summarized the average activity of these populations. The number of integration variables in the mean-field approximation is reduced to one for each neural population. Thus, it can be solved much more quickly and the parameter space can be scanned (Fig. 3.6A). Clearly, this obliged us to check the consistency of the mean-field calculations with the simulated
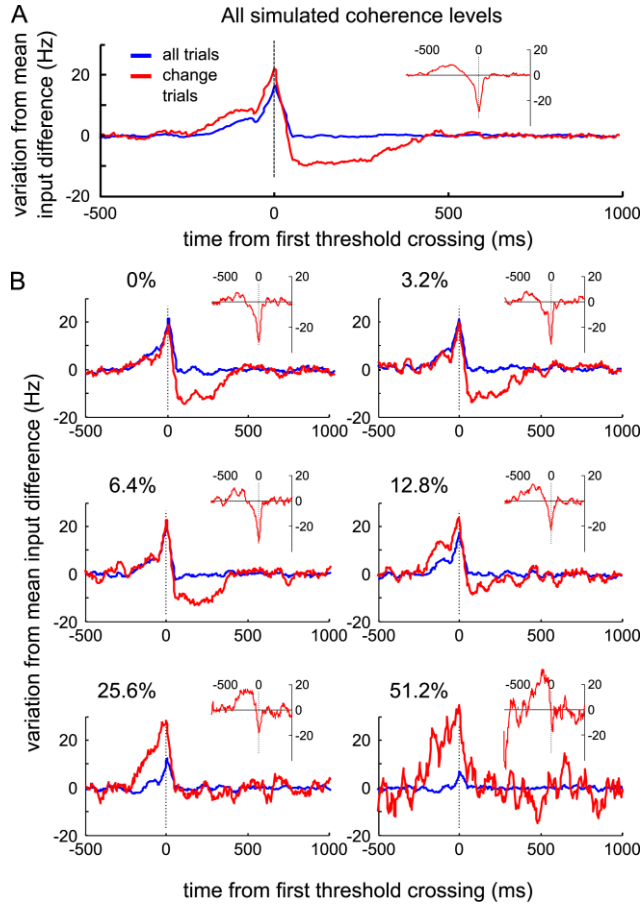
Fig. 3.5 Influence of input noise on changes of mind.

The variation from mean input difference of the selective populations, signed according to which pool first crossed the decision threshold, was averaged aligned to the first threshold crossing for all trials and all change trials. The insets show the input variation for change trials aligned to the second threshold crossing. (A) Mean across all coherence levels. (B) Separated by motion coherence.

activity of the full spiking network. We did this by running both sorts of simulations with the same parameters at key points in their parameter space (see 3.3.5).

By solving the mean-field equation for a set of initial conditions (here the initial firing rates of each neural population) one obtains the approximated average firing rate of each pool, when the system has settled into a stationary state. These stationary states correspond to the stable states or attractors of the system (Fig. 3.6A, thick black lines). The

unstable fixed points denote the border of the "basins of attraction" of the stable states (Fig. 3.6A, dotted black lines).

The present model has three qualitatively different dynamical regions across the range of symmetric inputs to the selective populations from 0 to 200 Hz. These dynamical regions are separated by fixed-point bifurcations, where a stable fixed point becomes unstable or vice versa. For small inputs the spontaneous state (↓↓), where both selective pools fire at low firing rates, is still stable (Fig. 3.6A, blue shaded region). At about 20 Hz the system crosses the first bifurcation and the spontaneous state becomes unstable. The network then operates in a region of categorical decision-making, where only the decision state (↓↑) is stable. There, one selective pool will settle at the upper branch, the other will decay to the lower one. With sufficiently high selective inputs (> 125 Hz) a symmetric "double-up" state becomes stable (↑↑), where both selective populations fire with intermediate, elevated rates (Fig. 3.6A, grey shaded region). Because of the strong recurrent connections within the selective populations, the decision state is stable over the whole range of inputs shown and the spontaneous- and symmetric-state bifurcations are "subcritical pitchfork bifurcations".

The above conclusions still hold if, instead of symmetric selective inputs as in Fig. 3.6A, biased inputs are applied, favoring one selective population against the other. In that case the double-up state still exists, but the pool with positive bias will fire at a higher rate than the one with negative bias. The higher the bias, the more will the firing rates of the two selective populations differ in the double-up state. In addition, the basin of attraction of the decision state grows for the favored population at the expense of the other, making wrong choices less likely (Wong and Wang, 2006; Wong et al., 2007).

The mean-field approximation in general provides an accurate qualitative picture of the attractor landscape. Nevertheless, also quantitative conclusions can be drawn from the analysis. However, there is typically a shift of the predicted fixed points in comparison to the attractors of the spiking network, mainly due to the additional finite size effects in the spiking network (Brunel and Wang, 2001; Marti et al., 2008). To obtain a measure for this discrepancy, we performed network simulations to determine the fixed points of the full spiking model for some discrete selective input amplitudes (see A.1.3), shown as blue crosses in Fig. 3.6A. At 150 Hz selective inputs the symmetric state was first found to be stable for more than 3,000 ms in 9 out of 100 trials. The real second bifurcation point of the spiking network is thus shifted by about 25 Hz to higher inputs (i.e. to the right) with respect to the mean-field predictions. The input amplitude of the spiking simulation for which changes of mind can be obtained with the attractor model (155 Hz) lies close to this second bifurcation point. Note that in the spiking simulation

the dip of firing activity at motion onset marks the start of the transition to the decision state. The initial firing rates of the selective populations (about 25-30 Hz) are therefore located close to the symmetric attractor.

As a consequence of the proximity to the symmetric attractor, the decision process is prolonged (Wong and Wang, 2006; Roxin and Ledberg, 2008), making changes of mind more probable. A change of
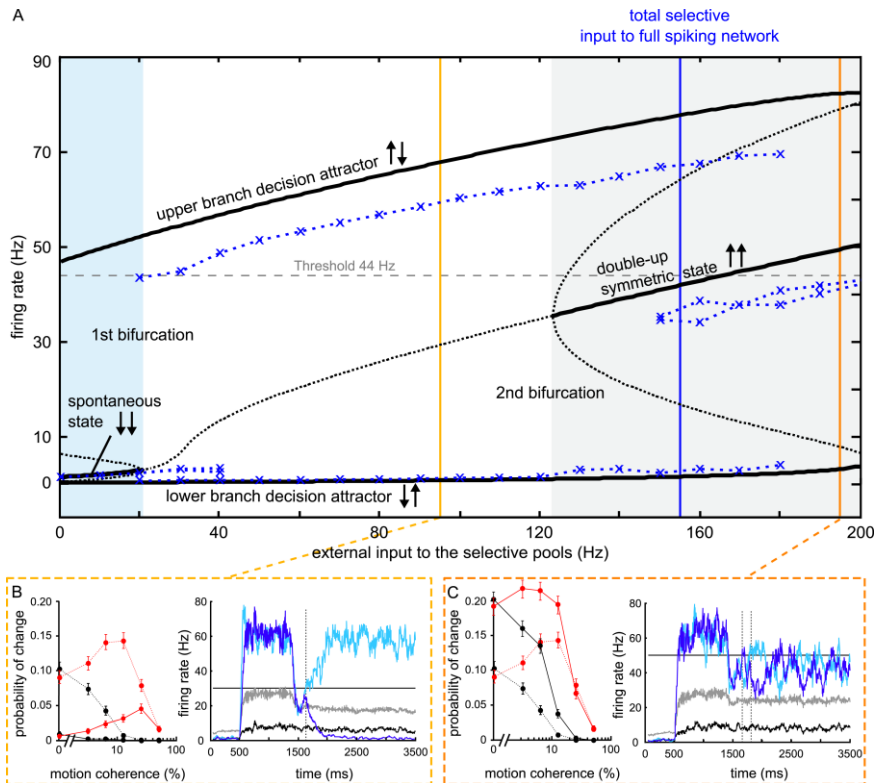
Fig. 3.6 Proximity to bifurcation is important to obtain changes of mind.
(A) Mean-field analysis of attractor network. For the parameters used in the spiking model simulation, the stable (solid black line) and unstable (dotted black line) fixed points were calculated with the mean-field approximation over a range of external inputs, applied symmetrically to both selective pools (0% coherence) from 0 to 200 Hz in steps of 1 Hz, in addition to the background input of 2.4 kHz to all neurons. The blue crosses show the fixed points of the spiking-neuron model for several discrete selective input amplitudes. (B, C) Changes of mind and single trial examples for lower (B) and higher (C) network inputs (yellow and orange lines in (A)). All parameters and the motion input were the same as in the other simulations, only the target input after motion onset was set to 25 Hz for (B) and to 125 Hz for (C). Dashed lines in the left panels give changes of mind from Fig. 3.2A for comparison. Red: changes to correct, black: changes to wrong choice. Color of single trial firing rates are the same as in Fig. 3.3.

mind is possible until one pool crosses the unstable fixed point (Fig. 3.6, dotted line between symmetric state and the decision branches) and falls too deep into the basin of attraction of the decision state, where only strong input fluctuations can pull it out again. Note also that the double-up symmetric state lies below the decision threshold (44 Hz, horizontal dashed line) while the upper branch of the decision attractor ("winner") lies above. Taking the shift between the mean-field and spiking-network attractors into account, the decision threshold of 44 Hz coincides approximately with the unstable fixed point and thus with the border between the basins of attraction of the double-up and the decision state. A change of mind can consequently be interpreted as a transient that comes very close to or even surpasses the unstable fixed point, but, because of contrary evidence or fluctuations, does not escape towards the upper decision state and eventually loses the competition.

3.3.5 Verification of mean-field prediction by spiking simulations

Although the above-presented notion of changes of mind is consistent with the mean-field attractor picture, the accuracy of the approximation is known to be especially weak close to bifurcation points (Brunel and Wang, 2001; Marti et al., 2008). The mean-field conclusions on the frequency of changes of mind thus have to be validated by simulations with the full spiking network.

Therefore, we performed spiking simulations for all coherence levels for different selective inputs to further demonstrate the importance of the system's proximity to the symmetric-state bifurcation (Fig. 3.6B,C, yellow and orange lines in Fig. 3.6A). All network parameters and the motion input were kept identical to the simulations presented above. The selective inputs were changed by varying the target input after motion onset. The decision thresholds were adjusted so that the model with altered selective inputs fit the experimental reaction times and performances (Fig. 3.A.3). For 25 Hz target input (and thus a total selective input of 95 Hz at 0% motion coherence), considerably less changes of mind were obtained, especially for low motion strength, despite the low decision threshold of 30 Hz. By contrast, with a target input of 125 Hz the model predicted too many changes at low motion coherence. More importantly, in most of the low coherence trials with high target input the selective pools did not leave the symmetric state (Fig. 3.6C, Fig. 3.A.3B). Contrary to the concept of using the attractor states to determine the decision outcome, here, even large fluctuations do not necessarily lead to a transition towards the decision attractors. By contrast, close to the bifurcation point, fluctuations will eventually lead to an escape from the symmetric state.

These additional simulations also justify the use of t_{ND} as a time-out for changes: turning the motion stimulus off with movement initiation would correspond to stopping the motion input in the simulations at t_{ND} after the first decision. The remaining symmetric target input of 85 Hz would be even lower than the selective inputs in the 95 Hz simulations with symmetric inputs (Fig. 3.6B). Thus, even if changes of mind were possible after t_{ND} they would be very unlikely.

Apart from the input to the selective populations, changing other network parameters will affect the location of the bifurcations. The general shape of the attractor landscape, however, is robust to gradual parameter changes. For example increasing (decreasing) the inhibitory connectivity ω_I shifts the whole attractor landscape to the right (left), which has a similar effect as decreasing (increasing) the selective inputs (Fig. 3.6) and likewise leads to fewer (more) changes (Fig. 3.A.4 and 3.A.5). This further confirms the crucial role of the symmetric state bifurcation for changes of mind in the attractor network.

3.3.6 Model predictions on bidirectional random-dot motion

As shown above, the frequency of changes of mind, as well as the simulated reaction times and performance of the attractor model, depend on the amount of common external input applied to both selective populations (Fig. 3.6). In Fig. 3.7 we give a more detailed analysis of simulated behavior with respect to common and biased external inputs, if the decision threshold is fixed at the standard decision criteria (44 Hz, 10 Hz difference). More precisely, we performed additional network simulations starting from various levels of equal external baseline inputs to both selective pools, indicated by different colors in Fig. 3.7: from 120 Hz in steps of 8.75 Hz to 155 Hz (the standard input close to the second bifurcation, used above to model the experimental changes of mind). On top of that, we varied the bias between the selective populations, again in steps of 8.75 Hz from 0 to 43.75 Hz (abscissa). In this input scheme, the pink and red dots correspond (approximately) to the standard input parameters used above at 0% and 25.6% motion coherence (here actually 25%). Increasing the baseline inputs leads to faster reaction times and overall more changes. Performance is less affected, but still decreases uniformly regardless of input bias.

An experimental equivalent for higher inputs to both selective populations might be obtained by increasing the overall dot density, or, alternatively, with bidirectional random-dot motion, similar to the three-alternative experiment by Niwa and Ditterich (2008). Independent coherent motion in two opposed directions allows comparing differences in the total sensory input while keeping the bias fixed. As an example, in

the case of 10% dots moving to the right and 20% to the left, fewer changes, larger reaction times and higher performance would be expected than for 30% of dots to the right and 40% to the left. Such an experiment should generally help to distinguish the nonlinear attractor model from linear diffusion models as used by Resulaj et al. (2009), which implement the accumulation of evidence as a single decision variable, encoding only the difference in sensory evidence, but not the absolute value for each direction. Still, changes in the input variance might affect the diffusion model in a similar way as changes in the baseline input affect the attractor network (Fig. 3.8). Less variance in the input to the diffusion model leads to fewer changes, higher reaction time and better performance. Thus, to unambiguously distinguish the two types of models based on behavioral data, the experimental stimulus fluctuations should be controlled for.

Nevertheless, the two scenarios, input variation in the attractor model versus variance changes in the diffusion model, also differ in their predictions on the variance of the output firing rates across trials (compare Fig. 3.7D with Fig. 3.8D). While the variance across trials in the diffusion
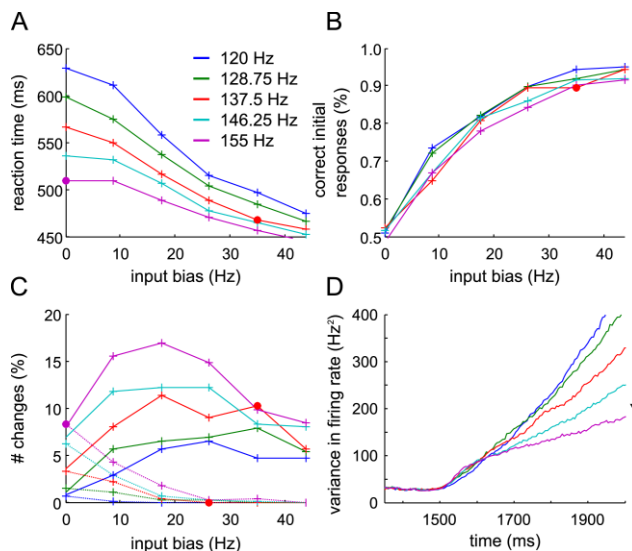
Fig. 3.7 Model predictions for different levels of common selective inputs.
The baseline external input, common to both selective populations, as well as the Mean reaction times (A), performance (B) and changes to correct (C, solid lines) and wrong (C, dashed line) alternative are plotted against the input bias between the selective populations. The decision threshold was fixed at the standard decision criteria (44 Hz, 10 Hz difference). 1,000 trials were simulated for each data point. (D) Evolution of the mean firing rate variance across trials for one selective population, starting from shortly before motion input onset (1,500 ms).

model intuitively increases with increasing input variance, in the attractor model it actually decreases with higher baseline inputs to the selective populations. The reason is again the vicinity to the second bifurcation, which impedes the escape to the decision attractors more the higher the inputs, leading to smaller variation in firing rate across trials. Neurophysiological recordings could thus distinguish the two mechanisms based on this higher order measure.

3.4 Discussion

Given the previous success of attractor models to simulate and explain behavioral and neurophysiological data of the RDM task (Wang, 2002; Wong et al., 2007; Albantakis and Deco, 2009) and decision-making in general (Deco and Rolls, 2006; Wang, 2008), here we made use of a binary attractor model with biophysically realistic neural dynamics to shed light on brain processes during changes of mind. We showed that, despite their fixed-point stability, attractor models are capable of capturing the essential aspects of changes of mind during the dynamic
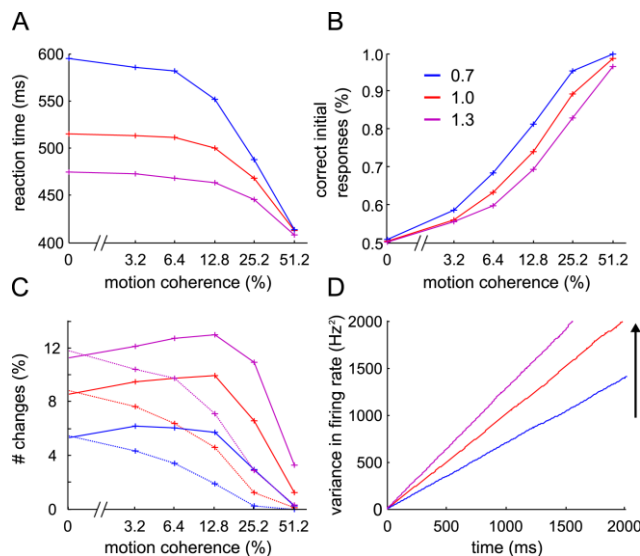
Fig. 3.8 Modifying the variance in the drift diffusion model.
Behavioral predictions of an extended linear accumulator-to-bound model, as used in Resulaj et al. (2009) for three different levels of input variance (0.7, 1.0, 1.3). Increasing the input variance leads to faster mean reaction times (A), worse performance (B) and more changes of mind (C). (C) Solid lines indicate changes to the correct alternative, dashed lines erroneous changes. 10,000 trials were simulated for each data point. (D) Evolution of the output variance with time.

transitions to the steady states. Moreover, a mean-field analysis revealed that the working point of the network, which fitted the experimentally observed changes of mind, is located close to a bifurcation, where a symmetric elevated state becomes stable. In the following we will discuss this and further model predictions on brain dynamics during changes of mind.

3.4.1 Distinction against alternative concepts for changes of mind

The presented attractor model offers a simple, yet biologically detailed, explanation for changes of mind with predictions on physiological recordings and the dynamical state of the brain region involved in the decision-making process. As in the bounded-accumulation model of Resulaj et al. (2009), a threshold crossing determines the initial choice, which can then be reversed by further processing of the remaining available information. Importantly, the linear accumulator model is not a reduced one-dimensional version of the attractor model. The mechanism behind the changes of mind is quite different. The attractor model is highly nonlinear: once the transient falls into the basin of attraction of the decision state, it is captured by the attractor and a change of mind is no longer possible, except for very strong fluctuations.

a) Comparison with previous studies of the attractor model

The original publication by X.J. Wang (2002) discussed decision reversal in the attractor model due to signal reversal, i.e. by explicitly inverting the motion input to the network. Similarly, Wong et al. (2007) studied the model behavior if short (100 ms) motion pulses were applied to the selective populations enhancing or weakening the coherent motion. There are two crucial differences between the "changes of mind" observed by Resulaj et al. (2009), which we dealt with in this chapter, and the previous approaches on "choice reversal": first, changes of mind here arise without explicitly inverting the motion evidence, solely by noise fluctuations in the RDM stimulus or, for the simulations, in the external selective input. Second, the inverted inputs in Wang (2002) and Wong et al. (2007) acted mainly before the decision threshold was crossed a first time and thus affected primarily performance. For a "true" change of mind, i.e. a first decision with a subsequent second threshold crossing, reversing inputs had to surmount the initial motion coherence substantially (Wang, 2002). In the present study, the input fluctuations inducing changes of mind are of about the same size as the fluctuations preceding the first threshold crossing (Fig. 3.5). This can be explained by the proximity to the second bifurcation, which delays the ultimate transition to the decision attractors and allows for initial fluctuation in the

output firing rate. Changes of mind, without explicitly reversing the input to the selective populations, are therefore not self-evident in the attractor model and occur only rarely, except for the dynamical regime close to the second bifurcation.

b) Comparison with the diffusion model

In order to reduce the free parameters and for physiological considerations, we set the non-decision time t_{ND} as time-out for changing after the first decision and used the same threshold for the first choice and a change of mind. By contrast, Resulaj et al. (2009) imposed a second independent threshold and an adaptable time-out for changes to fit their experimental results with an extended diffusion model. Thereby, they could account well for the participants' behavior and the frequency of changes. The predictions on neural activity by the one-dimensional model are, however, quite limited. In turn, we did not attempt a perfect quantitative fit to the data, but provided a neurodynamical explanation for changes of mind, based on the shape of the attractor landscape, which is robust to gradual parameter changes. Still, the simulated behavior fits the experimental data well. The attractor model only predicts slightly less erroneous changes and, hence, a larger difference in performance with and without changes in comparison to the participants' behavior and the diffusion model. This minor discrepancy might be accounted for by modifying the implementation of motion coherence: for simplicity we modeled coherent motion with a balanced input bias that affects both selective populations equally and grows linearly with increasing coherence (Eq. 3.2). Nonetheless, an unbalanced more positive bias, or a nonlinear increase with coherence (initially less for low coherence and more for higher coherence levels) would be plausible alternatives (Britten et al., 1993) that could provide a closer fit to the experimental data, without changing any of the predictions or conclusions presented in this study.

Although the validity of the two models cannot be distinguished based on their fits to the behavioral data of Resulaj et al. (2009), a slightly modified version of the RDM task with independent coherent motion in two opposed directions (Niwa and Ditterich, 2008), which allows comparing differences in the total sensory input while keeping the difficulty fixed, might give more information in that regard. The proposed attractor model predicts that the frequency of changing increases with higher sensory evidence for both alternative directions.

Apart from that, both of the above models assume that the brain continues to process incoming information after the initial decision. This hypothesis still needs to be verified by electrophysiological recordings. Another plausible mechanism is a reset of neural activity after the first threshold crossing. In the attractor model that would cause more changes

of mind. This can be understood easily for the 0% motion coherence case: a reset there means starting the decision process from scratch with again equal probability for both choices, while, in order to change decision for continuous processing, the transient first has to escape from the initial attractor. Moreover, resetting neural activity necessarily involves further mechanisms from external brain regions. In this study, however, we aimed to explain the changes of mind as an intrinsic feature of the decision-making process, based on nonlinear evidence accumulation with typical noise fluctuations.

3.4.2 Two mechanisms for speed emphasis to obtain changes of mind

One requirement for intrinsic changes of mind in the attractor model is a relatively low (first) decision threshold. A low threshold implies fast reaction times and comparatively low performance and thus corresponds to an emphasis on speed against accuracy (Ratcliff and Smith, 2004; Palmer et al., 2005; Lo and Wang, 2006). Indeed, Resulaj et al. (2009) suggest that time pressure induces changes of mind, as fewer changes were observed when participants were instructed to perform more slowly. Moreover, a low threshold in the attractor model leads to the experimental prediction of a bimodal build-up of the mean firing rates (Fig. 3.4C). After an initial uniform ramping activity that terminates already close to the threshold, the slopes of the average firing rates diverge rapidly for the various motion coherences. As coherence-dependent differences in mean ramping activity only set in near the decision threshold, differences in reaction time with motion strength are rather small. The reaction times of the three participants from Resulaj's experiments are in fact very fast and differ by less than 150 ms between 0% and 51.2% motion strength in comparison to over 400 ms in previous studies with well-trained monkeys (Roitman and Shadlen, 2002) or human subjects without explicit instructions on speed or accuracy (Palmer et al., 2005). More generally, neurophysiological recordings along the lines of our predictions in Fig. 3.4F could yield further experimental evidence on the existence and value of an absolute decision threshold in LIP.

Apart from the decision boundaries, the speed-accuracy tradeoff can, theoretically, be controlled by a second mechanism: Roxin and Ledberg (2008) showed that, in a reduction of the attractor model to a one-dimensional nonlinear diffusion equation, higher common inputs to both selective populations lead to a decrease in performance and reaction times (Fig. 3.7)[9]. Supporting experimental evidence comes from several recent

---

[9] Strictly speaking, the one-dimensional reduction, and hence also the monotonic dependence of the mean input to speed and accuracy, are analytically only valid

fMRI studies, where an increase in the activity of neural integrators was observed with speed emphasis (reviewed in Bogacz et al., 2010). The mean-field analysis and complementary simulations with different selective inputs (Fig. 3.6) revealed that, in order to explain the frequency of changes found by Resulaj et al. (2009), high common inputs to the selective pools are required in addition to a low threshold. Therefore, we suggest that, physiologically, both mechanisms to implement a speed emphasis are essential to explain the experimentally observed changes of mind: high selective inputs and a low decision threshold.

3.4.3 Physiological relevance of the bifurcation between decision-making and double-up state

Previous analyses of the binary attractor model for decision-making (Wang, 2002; Wong and Wang, 2006; Marti et al., 2008) all focused on a region in the vicinity of the first bifurcation, where the spontaneous state becomes unstable. There, performance is high and reaction times are rather long, because of long stimulus-integration times. Recently, also the "double-up" symmetric state gained relevance in connection with target presentation (Wong et al., 2007; Furman and Wang, 2008; Albantakis and Deco, 2009), since consistent experimental evidence was found for high firing rates just before stimulus presentation (Roitman and Shadlen, 2002; Huk and Shadlen, 2005; Churchland et al., 2008; Kiani and Shadlen, 2009). Assuming high selective inputs with target onset, the double-up state can explain neural activity prior to the decision-making period. Furthermore, in (Soltani and Wang, 2009) cue inputs that arrive while the system is in the symmetric up-state add up to determine the network's starting point for subsequent decision-making, thereby implementing probabilistic inference.

If neural activity in decision-related areas actually evolves according to an attractor landscape, as proposed by this and previous studies (reviewed in Wang, 2008), the dynamical system has to cross a bifurcation in order to switch between the double-up state, effective during target presentation, and the decision-making regime, during random-dot motion. Yet, experimental indications that would suggest any physiological relevance of this second bifurcation for brain dynamics during decision-making have been lacking.

---

close to the first network bifurcation, where the spontaneous symmetric state becomes unstable (Roxin and Ledberg, 2009). In the presented model, the optimal working point of the system in order to account for the experimental data, however, lies close to the other bifurcation, where the symmetric state reappears with elevated firing rates in both selective pools (Fig. 3.6A). Nevertheless, the mean-field analysis and complementary simulations showed that the monotonic speed-accuracy relation to the selective inputs apparently still holds.

In this study, we found that the attractor model best captures the behavioral data and changes of mind observed in the experiments of Resulaj et al. (2009), if the system lies in the proximity of the second bifurcation. We thus proved that all input regimes of the binary attractor model are consistent with particular aspects of the decision-making process and thereby confirmed the suitability of the attractor model to describe neural dynamics. Consequently, we predict that the brain operates over the whole range of inputs that enable decision-making, dependent on the pressure for speed or accuracy, instead of switching between two discrete input levels for decision-making and target representation. This could be tested pharmacologically by gradually blocking inhibition in the decision-related brain areas: decreasing inhibition shifts the working point of the system closer to the bifurcation (Fig. 3.A.5). Thus, decreasing reaction times, lower accuracy and more changes would be expected, until the double-up symmetric state becomes stable, where decision-making might consequently be impaired completely for low coherence levels.

Taken together, we showed that changes of mind arise naturally in an attractor model of perceptual decision-making by emphasizing reaction speed against accuracy. We suggest that this speed-accuracy tradeoff is physiologically implemented by both, threshold adaptation and increasing symmetric inputs. Moreover, we found evidence for the physiological relevance of a so far unregarded bifurcation in the binary attractor model and thereby confirmed the general accordance of attractor networks with neural processes. Finally, we provided predictions on a new experimental paradigm, which might help to distinguish between nonlinear attractor and linear diffusion models.

3.A Chapter appendix

3.A.1 Network simulations without target stimulus

To investigate the network's behavior without a time-dependent target signal, we performed additional simulations with the same network parameters as in the main text of Chapter 3, but without the initial phase of the target signal (1,000 trials for each motion coherence). Thus, at 500 ms an input of 85 Hz was applied together with the motion signal (70 Hz ± bias). The network is nevertheless capable of decision-making, although with about 50 ms larger reaction times (Fig. 3.A.1A). The first-choice performance is unaffected by the missing initial target signal. Changes of mind still occur, albeit fewer, which leads to less performance improvement with changing (Fig. 3.A.1B).

The slightly longer reaction times and fewer changes are expected without the high initial target input, as in that case the firing rate transients evolve from the spontaneous state with very low firing rates in both selective populations. Therefore, the transients are more distant to the decision state and double-up state at the start of the motion signal. The match to the experimentally observed reaction times and changes of mind might be improved with threshold adaptation. However, if no targets are presented during the task, the subjects' behavior is likely to differ as well to some extent. Nevertheless, the fact that, even without the initial target signal, changes of mind still occur repeatedly in the model, due to the proximity of the second bifurcation, further strengthens the robustness of our findings.
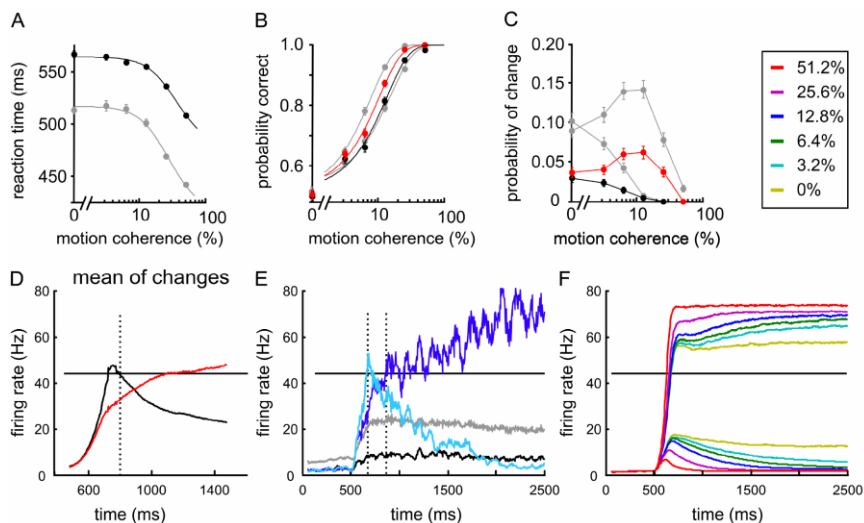
Fig. 3.A.1 Simulations without target stimulus.
(A) Reaction time. (B) Initial (black) and final (red) performance. (C) Changes of mind. (A-C) Results from the main text are shown in gray for comparison. (D) Mean firing rate of changes of mind trials aligned to first threshold crossing. (E) Single 0%-motion coherence trial with change of mind. (F) Mean firing rates of all motion coherences (color legend displayed above). Colors as in Fig. 3.2 and Fig. 3.4.

3.A.2 Robustness of simulation results to variation in decision parameters

The influence of the difference criterion on the variables of interest, reaction time, performance and changes of mind, is shown in Fig. 3.A.2 A-C. With respect to the standard decision and change criteria (a 44 Hz threshold and a difference of 10 Hz between the firing rates of the selective populations), a smaller difference criterion (5 Hz) (but the same threshold) leads to faster reaction times and somewhat more changes. A larger difference has the opposite effect. If, in addition to a smaller difference, the threshold is slightly adapted to 45.5 Hz, the results from the standard criteria are fully recovered. With respect to the time distribution of the changes of mind, the difference criterion affects mostly the amount of early changes, if the threshold is not adapted (Fig. 3.A.2D). In Fig. 3.A.2E an example trial is displayed where, during the initial joint
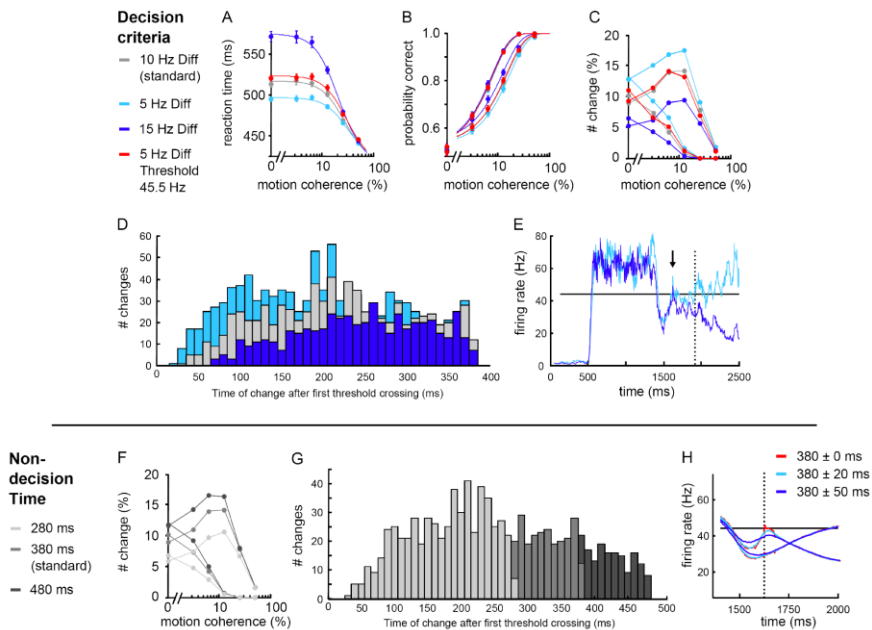
Fig. 3.A.2 Robustness of simulation results to variation in decision parameters.
(A-E) Influence of the difference criterion. (F-H) Variability in the non-decision time. (H) Examples of simulated mean activity for change trials as in Fig. 3.4F. Here, however, the non-decision time in each trial is drawn randomly from a normal distribution with mean 380 ms and a standard deviation of 20 ms (light blue) or 50 ms (dark blue). The trials were aligned to first threshold crossing plus the deviation from the mean non-decision time.

increase of activity in the selective populations, both pools briefly cross the decision threshold. Due to the difference criterion this first joint crossing is not taken into account.

Varying the non-decision time mainly affects changes at low coherences levels. Even if the non-decision time in each trial is drawn randomly from a normal distribution, the switch in firing rates between the two selective populations should still be observable experimentally (Fig. 3.A.2H). Deducing a quantitative value for the decision threshold is however somewhat impeded by a broad distribution of non-decision times.

3.A.3 Varying the selective inputs

Keeping all other network parameters fixed, the network's working point depends on the mean inputs to the selective populations. Nevertheless, the attractor model can theoretically fit experimental reaction times and performance for a large range of selective inputs, if the threshold is adapted accordingly (Fig. 3.A.3).

With a target input of 25 Hz after motion onset, the decision threshold has to be lowered to 30 Hz to return approximately the same reaction times and performance as for the standard simulations with 155 Hz of selective inputs and a decision threshold of 44 Hz (Fig. 3.A.3A, grey lines in left and middle panel). As there are only very few changes, the performance with and without changes is very similar. For 125 Hz target
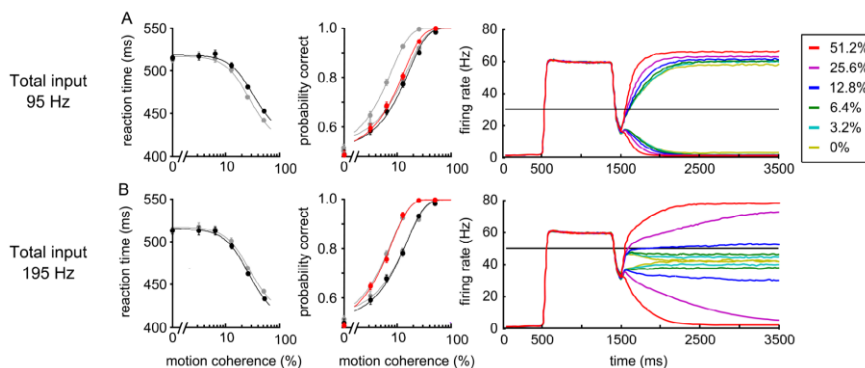
Fig. 3.A.3 Reaction times, performance and mean firing rate for different selective inputs.
(A) Reduced target input of 25 Hz and thus 95 Hz of total selective inputs with a decision threshold of 30 Hz. (B) Increased target input of 125 Hz and thus 195 Hz of total selective inputs with a decision threshold of 50 Hz. Colors as in Fig. 3.2 and Fig. 3.4.

input, the threshold has to be increased to 50 Hz to match the values obtained with 85 Hz target input (155 Hz total selective input). Note that the mean firing rates of correct trials up to 6.4% of motion coherence do not show a build-up to the decision attractor (right panel). A threshold crossing therefore is mainly caused by large fluctuations around the symmetric state and depicts a decision based on evidence integration only for higher motion strength.

3.A.4 Varying inhibition

Increasing inhibition has similar effects as decreasing the selective input, because the attractor landscape is effectively shifted to the right, towards higher selective inputs. This sets the network input (155 Hz) to the left of the spiking network bifurcation point, equivalent to lower inputs with the original inhibitory connection weights (see Fig. 3.6A).

In the same way, decreasing inhibition has similar effects as increasing the selective input. The attractor landscape is effectively shifted to the left, towards lower selective inputs, which sets the network input (155 Hz) to the right of the spiking network bifurcation point, corresponding to higher inputs with the original inhibitory connection weights.

The threshold has to be adjusted to 38 Hz and 50 Hz, respectively, to obtain comparable reaction times and performance at first choice. The simulations from the main text with $\omega_I = 1.125$ are displayed in grey in Fig. 3.A.4 and 3.A.5 for comparison. As for lower (higher) inputs (Fig. 3.6 A and B) increasing (decreasing) the inhibitory connections leads to fewer (more) changes especially at low motion coherence.
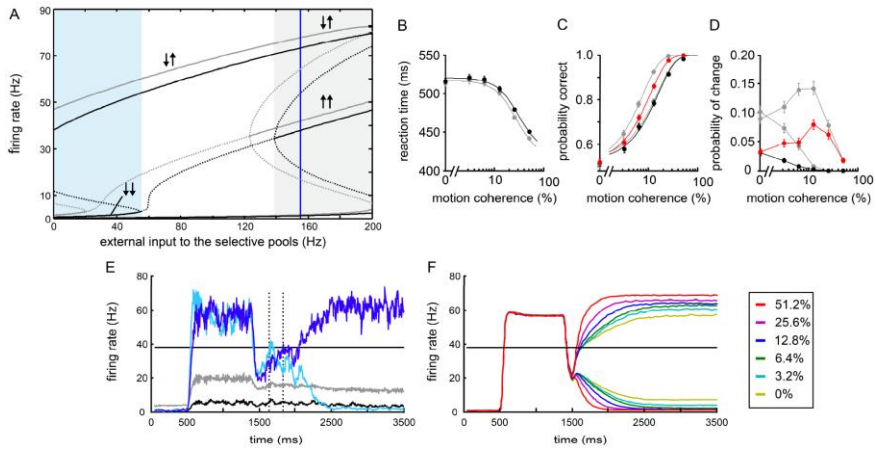
Fig. 3.A.4 Spiking simulation with increased inhibition.
(A) Mean-field analysis of network with higher inhibition ($\omega_I = 1.425$ $\omega_I = 0.825$ instead of 1.125, all other parameters and inputs as before). (B-D) Simulated behavior. (E) Single trials with 0% motion coherence and change of mind. (F) Mean firing rates for correct first choices. Colors as in Fig. 3.2 and Fig. 3.4.
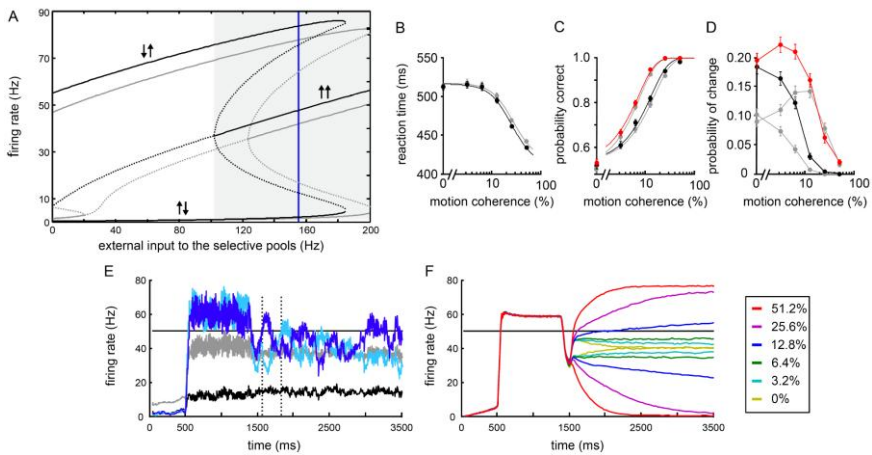
Fig. 3.A.5 Spiking simulation with decreased inhibition.
(A) Mean-field analysis of network with lower inhibition ($\omega_I = 0.825$ instead of 1.125, all other parameters and inputs as before). (B-D) Simulated behavior. (E) Single trials with 0% motion coherence and change of mind. (F) Mean firing rates for correct first choices. Colors as in Fig. 3.2 and Fig. 3.4.

64

4 THE ENCODING OF ALTERNATIVES IN MULTIPLE-CHOICE DECISION-MAKING

The work presented in this chapter is published in PNAS (Proceedings of the National Academy of Sciences of the United States of America)[10]. *An abstract was published at the* CNS 2009 *conference*[11].

4.1 Introduction

In the last chapter, we have been concerned with the applicability of attractor-networks to the process of choice-reevaluation and thus focused on a temporal extension of binary perceptual decision-making to the time after the initial choice. Here, we take another step and address the extension from binary to multiple-choice decision-making.

Already decades ago decision-making between multiple alternatives was the subject of psycho-physical reaction-time studies, which revealed an increase in reaction times with the number of choices (Hick, 1952). With the objective of shedding light upon the neural mechanisms that underlie decision-making, experimental and theoretical studies subsequently focused on the simplest case of binary choices (Smith and Ratcliff, 2004; Gold and Shadlen, 2007; Wang, 2008).

This led to the identification of the lateral intraparietal area (LIP) as a candidate for bounded integration in the decision process. As stated above, the neural activity in this cortical area correlates with the choices and reaction times of monkeys performing the random-dot motion (RDM) task (Fig. 4.1A), the established paradigm to test for accumulation and integration of evidence during decision-making (Shadlen and Newsome, 1996, 2001; Roitman and Shadlen, 2002).

The biophysically realistic spiking-neuron model of LIP proposed by Wang (2002) successfully simulated behavioral and physiological data from the binary RDM task. It is based on attractor dynamics and winner-take-all competition of two discrete selective populations of neurons (pools), each representing one alternative. In this chapter, we propose a possible extension of the binary attractor model to multiple alternatives.

Several basic connectionist models of neural networks already addressed choice behavior regardless of the number of alternatives (Usher

---

[10] Albantakis L, Deco G (2009) The encoding of alternatives in multiple-choice decision making. *Proc Natl Acad Sci U S A* 106: 10308–10313.

[11] Albantakis L, Deco G (2009) The encoding of alternatives in multiple-choice decision making. Eighteenth Annual Computational Neuroscience Meeting: CNS*2009 Berlin - BMC Neuroscience 10(Suppl 1):P166.

and McClelland, 2001; McMillen and Holmes, 2006; Bogacz et al., 2007; Niwa and Ditterich, 2008). Yet, it is difficult to relate these rate models to explicit physiological processes during multiple-choice decision-making.

Only recently, experimental studies extended the RDM paradigm to more than two alternatives (Churchland et al., 2008; Niwa and Ditterich, 2008). In particular, Churchland et al. (2008) compared behavioral data and recordings from single LIP neurons of macaque monkeys performing a 4-choice RDM task with the original 2-alternative task. Reaction times and error rates for four alternatives were found to be longer and higher, respectively, consistent with earlier studies (Hick, 1952). In an additional control condition with two response targets (R-targets) separated by 90º (90º-case), monkeys needed longer to decide than in the standard (180º) 2-alternative case, but performed with the same accuracy. Notably, the experiments of Churchland et al. (2008) provided the first electro-physiological data on a 4-alternative decision task.

Two theoretical studies (Beck et al., 2008; Furman and Wang, 2008) subsequently proposed continuous models of multiple-choice decision-making. Both models can account for important findings of Churchland et al. (2008). One, by Beck et al. (2008), focused on the implementation of probability distributions and optimality, while the model of Furman and Wang (2008), like our model, features high biophysical detail. It combines Wang's (2002) discrete 2-alternative model with a hypercolumn model (Ben-Yishai et al., 1995), where a ring of neurons represents continuous directions of motion. However, it cannot account for the condition with two R-targets spaced 90º apart. What is more, their model requires regulatory mechanisms depending on the number of alternatives, like an adaptation of the target input and an external top-down control signal during the decision-making period.

Here, we propose a different approach to extend the biophysically-based binary decision model (Wang, 2002). Instead of a continuous representation we increased the number of discrete neural populations that encode the possible alternatives (Fig. 4.1B,C). Previously, networks with discrete populations have been adjusted to exhibit winner-take-all competition for one particular set of choice alternatives (Wang, 2002) or memory states (Brunel and Wang, 2001). In our study, we analyzed how the network's competition regimes could be brought into accord for different numbers of alternatives. With a common parameter set for the 2-, 4- and 90º-case, we successfully simulated all experimental paradigms tested by Churchland et al. (2008), without the need of any number-of-choice dependent mechanism. Besides, we found that encoding decision alternatives by populations of neurons with a big relative pool size, a high "coding level", favors a common decision regime for two and four choices. Taken together, our results indicate a physiological advantage of a pooled, multi-neuron representation of choice-alternatives.
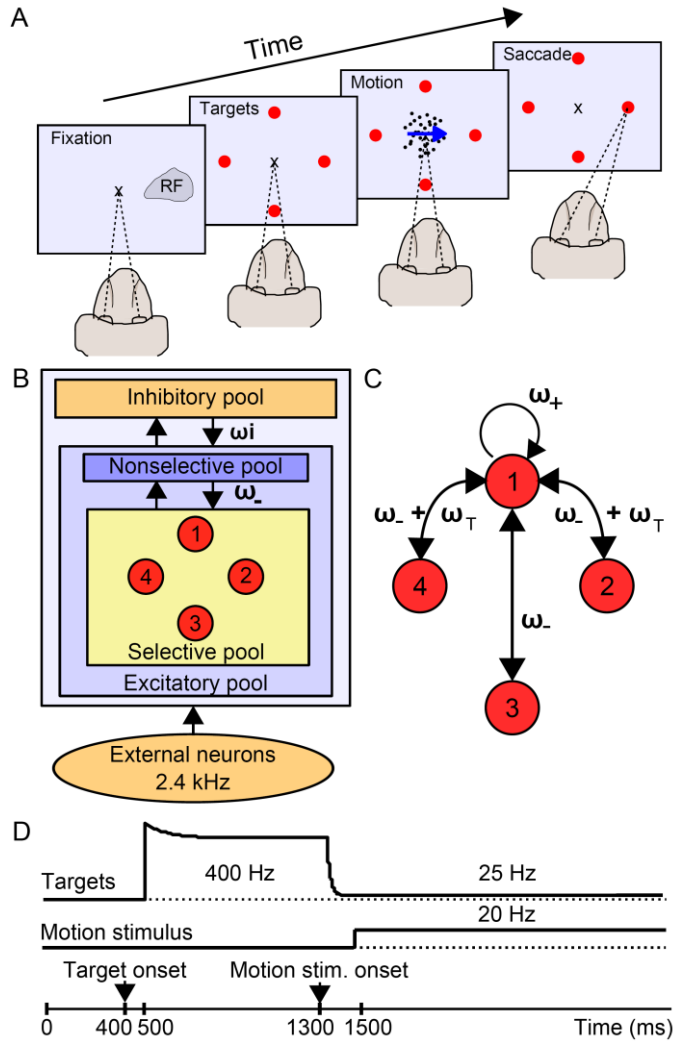
Fig. 4.1 Experimental design, network architecture, and stimulation protocol.
(A) The multiple-choice RDM task as in (Churchland et al., 2008). (B) Diagram of the spiking neural network model. (C) Connectivity between selective populations shown representatively for one pool. Synaptic weights to and from neighboring pools are additionally enhanced by a value ω_T. (D) Time course of input to the selective pools. Depending on the number of alternatives a target input was applied to either two or four pools.

Summarized, our main findings are:

- The discrete attractor model with spatially graded connectivity accounts for the all experimental conditions tested in a 2- and 4-choice RDM task (Churchland et al., 2008).

- No extra top-down regulation mechanisms are necessary to adapt the network to the number of choices.

- The network's capacity for choice-number independent decision-making is higher the more neurons encode each choice-alternative.

4.2 Methods

4.2.1 Experimental Paradigm

Churchland et al. (2008) tested two macaque monkeys on a multiple-choice version of the RDM task (Fig. 4.1A) and measured the primates' decision-behavior (reaction times and performance), as well as the decision-related single cell activity of LIP neurons.

As in the classic 2-alternative RDM task, the monkey initially fixated on a central spot, while the R-targets indicated the possible alternatives for the direction of coherent dot-motion (Fig. 4.1A, red dots). The R-targets continued to be present throughout the full trial. One of the R-targets was always located in the response field (RF) of the recorded LIP neuron. After a delay, a patch of dynamic random dots appeared, with a proportion of dots moving coherently toward one of the R-targets. The remaining dots were moving randomly. The amount of coherence controlled the task difficulty. The monkeys had to decide on the net direction of motion and to report their choices by a saccadic eye movement to the corresponding R-target.

Three experimental condition were compared, differing in the number and configuration of possible motion directions indicated by the R-targets: Either two opposing R-targets, four R-targets (90° apart) or two R-targets with an angular distance of 90° (90°-case) were presented to the monkey before the motion signal started. Given two alternatives, the monkeys either had to distinguish between opposing or perpendicular motion directions. The latter (90°-case) served as a control to distinguish the effects of different numbers of alternatives from the effects of a smaller angular distance.

4.2.2 Multi-alternative attractor network

a) Network structure and connectivity

As in the last chapter, the general network structure of our model for multiple-choice decision-making is based on the binary attractor network first presented by X.J. Wang (2002) (see 2.2.3). Nonetheless, two important modifications have to be noted: first, instead of just two selective populations, the excitatory neurons in the network are now subdivided into four selective populations, encoding the four possible motion directions (Fig. 4.1B,C, red). Each of the four selective pools contains $f \cdot N_E$ neurons, where f is termed the "coding level" of the selective pools and $N_E = 1,600$ the total number of excitatory neurons in the network. To fit the experimental data we used a coding level of $f = 0.2$. In order to analyze the effects of the size of the neural populations, which encode the possible choices, we varied the coding level in the mean-field analysis presented in Section 4.3.3.

As a second modification, we introduced a spatial connectivity component between the selective populations. As before, the recurrent connectivity ω_+ of neurons within a selective pool is higher than the connectivity ω_- between selective pools (Fig. 4.1C). On top of that, each selective pool in our model is thought to have two "neighboring" selective pools, corresponding to its two perpendicular R-targets, and one "opposing" pool, representing the R-target at 180º angular distance. We increased the connectivity between pools that represent neighboring R-targets by the weight ω_T, assuming a slightly higher correlation between the neighboring pools than between the anticorrelated opposing pools (Fig. 4.1C). Modeling the circular spatial distribution of the R-targets with ω_T allowed us to account for the observed experimental differences between the standard (180º) 2-choice condition and the 90º-case (see 4.3). Without ω_T, the model would not be able to distinguish between the two conditions.

Finally, a (fifth) nonselective pool of excitatory neurons emulates activity in the surrounding brain areas and a homogeneous pool of inhibitory neurons, connected to all excitatory neurons in the network, mediates global inhibition. For a detailed description of the network connectivity, dynamics and parameters please refer to the Appendix, in particular A.1.5.

b) Network inputs

As in Chapter 3, all neurons in the network receive an excitatory stochastic background input with a rate of 2,4 kHz. Moreover, in the present biophysically realistic attractor model, the four selective pools are thought to represent the populations of neurons in LIP where the spatial information about one respective R-target and the motion directed towards

this R-target are combined. Accordingly, we modeled the R-target and motion stimuli presented to the monkeys as shown in Fig. 4.1D.

Target input. Depending on the number and location of the R-targets in the different experimental conditions, the neurons of either all four selective pools, the two opposing pools or, in the 90°-case, two neighboring pools receive the same target input during the model simulation. We assume that the target input is passed on to the respective pools as a Poisson spike train with a time dependent firing rate of:

$$\nu_{\text{target}} = \begin{cases} 0 \text{ Hz} & 0 < t < t_{\text{target}} + 100 \text{ ms} \\ (400 + 100\,exp(-(t - t_{\text{target}} - 100 \text{ ms})/\tau_1))\ \text{Hz} & t_{\text{target}} + 100 \text{ ms} \leq t < t_{\text{motion}} + 80 \text{ ms} \\ (25 + 375\,exp(-(t - t_{\text{motion}} - 80 \text{ ms})/\tau_2))\ \text{Hz} & t \geq t_{\text{motion}} + 80 \text{ ms} \end{cases}$$

$$(4.1)$$

where t_{target} = 400 ms and t_{motion} = 1,300 ms denote the onset times of the target and the motion stimulus, respectively. The initial exponential decay, τ_1 = 100 ms can be explained by short term adaptation. The dip in the firing rate is modeled by an exponential decrease of the target input with τ_2 = 15 ms, starting with a latency of 80 ms after motion stimulus onset (Wong et al., 2007; Furman and Wang, 2008).

Motion input. The motion stimulus without coherent motion (0% coherence) was simulated as a Poisson spike train with a time invariant rate of ν_{motion} = 20 Hz to all selective pools starting at 1,500 ms. Coherent motion was modeled as a positive bias to one selective pool, balanced by a reduction of the motion input in the other three selective pools to keep the total motion input to the network constant. A motion coherence of 100% thus corresponds to a bias of 60 Hz to one selective pool, resulting in a motion input of 80 Hz to this particular pool and 0 Hz to the other selective pools. We simulated 10 motion coherences: 0%, 1.67%, 3.33%, 5%, 8.33%, 12.5%, 25%, 50%, 75% and 100%.

Importantly, the motion input is received by all selective neurons in the network, whereas the target input is just applied to the particular pools corresponding to the possible choices. Again, this means that for four alternatives all four selective pools receive the target input, for the standard 2-choice condition pools 1 and 3 receive the target input, and for the 2-choice 90°-case the target input is applied to selective pools 1 and 2. Therefore, the target input is what distinguishes the different experimental conditions in the model simulations.

Our network is generally capable of decision-making even without the initial target signal before the onset of the motion input (Fig. 4.A.1). If the external input lies within the range of decision-making, the shape of the target input affects reaction times, but not the network's capacity of decision-making.

c) Decision threshold and network simulations

A decision is reached in the model when one selective pool crosses a threshold of 50 Hz and surpasses the other selective pools by at least 5 Hz. The reaction time (RT) is calculated as the time difference between the motion input onset (1,500 ms) and the time the threshold is reached, plus an additional non-decision time $t_{ND} = 280$ ms to account for saccade initiation and execution, and the signal latency of the RDM stimulus to arrive in LIP (Roitman and Shadlen, 2002; Churchland et al., 2008).

4.3 Results

4.3.1 Comparison to behavioral data

Churchland et al. (2008) measured the accuracy and speed of the monkeys' choices for several motion coherences (Fig. 4.2B, D). Fig. 4.2A and C show the reaction times and performance, i.e., the fraction of correct choices obtained by the model simulations.

In the experiments and also in the simulations the RTs were longer for four possible alternatives than for two[12]. RTs in the 2-alternative 90°-case were intermediate, with larger differences at lower motion strengths.

Starting at chance level (50% for two and 25% for four alternatives), the accuracy increased until it reached 100% for high motion strengths (Fig. 4.2C,D). Except for very high coherences, choices between four alternatives were less accurate than binary decisions, also in comparison with the 90°-case. There, the monkeys performed as well as in the standard two-choice case. In the simulations, the accuracy in the 90°-case resembles the standard binary case (Fig. 4.2C), with somewhat higher values at intermediate motion coherence.

In summary, although we did not attempt a perfect quantitative fit to the experimental data, the psychometric functions obtained by our model simulations match the experimental observations very well in all relevant aspects.

4.3.2 Comparison to neurophysiological recordings

Behavioral differences between the 2- and 4-alternative and the 90°-case must be based on differences in the temporal evolution of the firing rates during the decision process. The proposed attractor network can generally be viewed as representing a local microcircuit in area LIP. In

---

[12] At low motion coherence some simulated trials had to be excluded in the four-alternative condition (at most 2%) and in the 90°-case (at most 5%), because they failed to reach a decision within the simulation time of 4,000 ms.
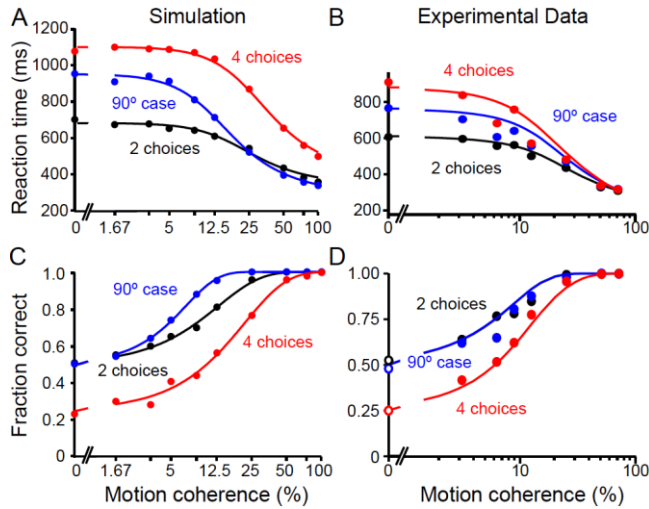
Fig. 4.2 Speed and accuracy of simulated decisions and comparison to experimental data.
(A) Simulated mean RTs of correct trials as a function of motion coherence, fitted by a hyperbolic tangent function (Eq. A.35). (B) Mean reaction times of two monkeys performing the RDM task. (C) Simulated psychometric functions, fitted by Weibull functions (Eq. A.36). 1,000 trials were simulated for each data point. (D) Experimentally observed performance of monkeys in the RDM task. (B, D) Adapted with permission from Churchland et al. (2008). Please note that the scales of the simulated and experimental data are identical. For better assignment in (A, C) the simulated motion coherence values are used as labels.

Fig. 4.3 the temporal evolution of the simulated firing rates is displayed for single trials and trial averages at zero motion coherence for each condition (2-, 4- and 90º-case).

The stochastic nature of the network due to finite-size effects and Poisson inputs allows decision formation even for unbiased inputs (Deco and Romo, 2008). The spike raster-plot (Fig. 4.3A) displays a four-choice trial with 0% motion coherence where all selective neurons receive the same input (10 sample neurons are shown from each selective pool, labeled 21 to 60). Eventually (at about 2,000 ms) the symmetry is broken and, in this case, the activity of the left pool (purple) shifts to an attractor state with increased activity (up-state).

The time course of the simulated neural activity is in good agreement with the experimental observations of LIP neurons (Huk and Shadlen, 2005; Churchland et al., 2008; Kiani et al., 2008): Throughout the target period, between target input and motion stimulus onset, the selective pools representing the possible direction alternatives exhibit elevated firing rates, followed by a "dip" to lower activity after motion onset and

subsequent ramping activity. With the arrival of the motion signal the integration process starts and a decision is finally made, characterized by an activity build-up in the winning pool.
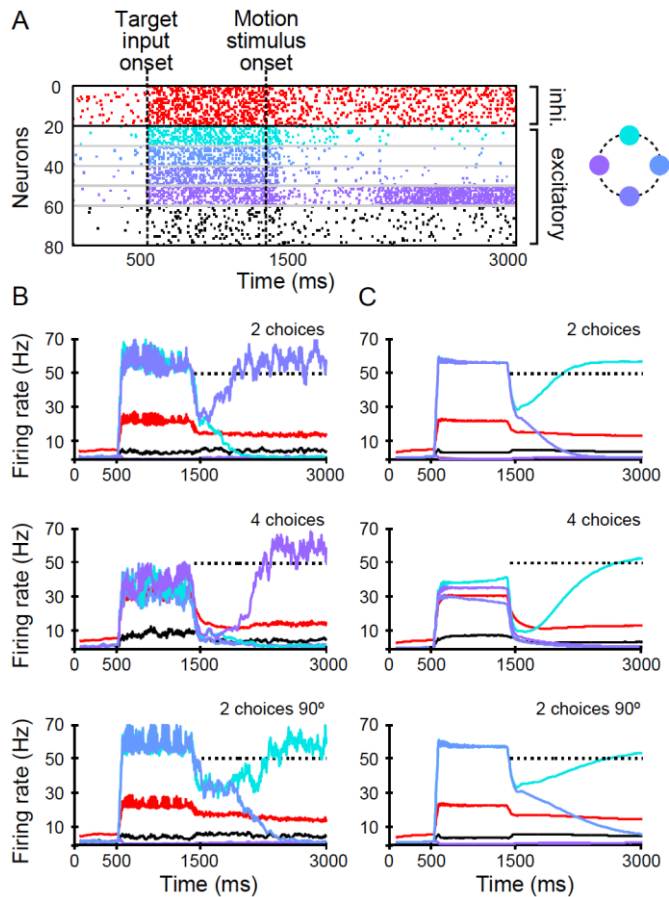
Fig. 4.3 Simulated temporal evolution of firing rates at 0% motion coherence.

(A) Spike raster plot for sample neurons of each network pool in a four-choice trial (the same as shown in B, middle panel): 20 inhibitory neurons (red, top), 20 nonselective neurons (black, bottom) and 10 neurons of each selective pool colored according to the schematic illustration of the R-target locations (right). (B, C) Population averaged firing rates for single trials (B) and trial average over 1,000 network simulations (C). Red and black lines denote inhibitory and nonselective pools. Selective pools for single trials are colored as neurons in (A). For the trial average, the "winning" neural pools were averaged (cyan). Respectively, the "loosing" pools were averaged according to their inputs and relative R-target location with respect to the winning population.

In the model simulations, the initial dip in the firing rate is caused by a reduction of the target input with an assumed latency of only 80 ms before the motion signal is supposed to arrive in LIP with a latency of 200 ms (Eq. 4.1). Possible physiological origins for the dip in neural activity are divided attention or upstream inhibition of the target signal caused by the onset of the random-dot motion (Wong et al., 2007; Furman and Wang, 2008).

A major discovery of Churchland et al. (2008) was that, regardless of the number of R-targets and motion coherence, the decision process is terminated at one stereotyped activity threshold. Differences between the 2- and 4-choice cases were instead observed during the target phase and in the early motion epoch. For four choices, the target response was on average 16.1 ± 1.6 Hz lower than for two choices, a difference that largely remained during the dip in firing rate at the onset of the motion stimulus (Churchland et al., 2008).

The attractor model matches these findings well, even quantitatively. The average firing rates during the target phase for four alternatives are about 20 Hz lower than for two possible choices (Fig. 4.3B,C, 57 Hz in the 2-choice and 90°-case compared to 36 Hz in the 4-choice case). As all network parameters are identical for the different task conditions, the differences in firing rate are caused solely by the shared feedback inhibition. The population activity of the inhibitory neurons, averaged over the time interval from 800 to 1,300 ms, is 32 Hz in the 4-choice condition compared to 22 Hz for the 2-choice and 90°-case. This can be explained by the higher total excitatory inputs from the selective pools for four alternatives: there, the target input is applied to all four selective pools, while for the 2-alternative case only two selective pools receive the target signal (pool 1 and 3 from Fig. 4.1C in the standard 2-choice case). In turn, the interneurons inhibit the activity of the pyramidal neurons more for four alternatives. As in the experiments, the conditional differences in firing rates persist during the initial dip after motion onset. Hence, the accumulation of evidence starts at lower values for four alternatives.

As can be seen in the trial average (Fig. 4.3C), the slope of the ramping activity for correct trials is similar for the 2- and 4-choice conditions. In the simulations, in accordance with the experimental findings (Churchland et al., 2008), a decision is reached when the population activity of one selective pool crosses a threshold of 50 Hz. The longer reaction times in the 4-choice task are therefore explained by the larger excursion of the neural activity to a common threshold, both, in experiment and simulations.

The 90°-condition is displayed in the lower panels of Fig. 4.3B,C. Firing rates during the target phase are similar to the standard two-choice case and slightly higher during the dip after motion onset. Interestingly, in the 90°-case the simulated average build-up activity is less steep than in

the other two conditions (Fig. 4.3C, lower panel). Therefore, longer RTs in the 90º-case must have a different cause than in the 4-choice case: as shown in the lower panel of Fig. 4.3B the initial symmetric state with high firing rates in both selective pools at ~1,500 ms is prolonged in the 90º-case compared to the standard 2-choice case. The reason is a stronger attraction of the symmetric state in the 90º-case due to the slightly higher connectivity ω_T between neighboring selective pools. This impedes competition between the selective populations and consequently leads to a smaller slope of average ramping activity towards the threshold and thus to longer reaction times. A shallower slope in the 90º-case was also observed experimentally (Churchland et al., 2008). The prolonged symmetric state, however, would hardly be measurable as the effect is lost in the trial average.

Moreover, these theoretical findings on conditional differences in the build-up activity without coherent motion further extend to biased motion coherence (Fig. 4.4). For all experimental conditions, the population averaged activity build-up is steeper with increasing motion coherence, which corresponds to decreasing task difficulty. This explains faster reaction times with higher coherence, as we have noted in the reaction time curves in Fig. 4.2A. Yet, even for higher levels of motion coherence, the ramping slope in the 90º-case remains less steep compared to the standard 2- and 4-choice cases, which have similar slopes.

Besides, the attractor model is capable of persistent activity due to the strength of its recurrent connections (Fig. 4.A.2), consistent with LIP neurons during working memory tasks (Shadlen and Newsome, 2001; Churchland et al., 2008; Kiani et al., 2008).
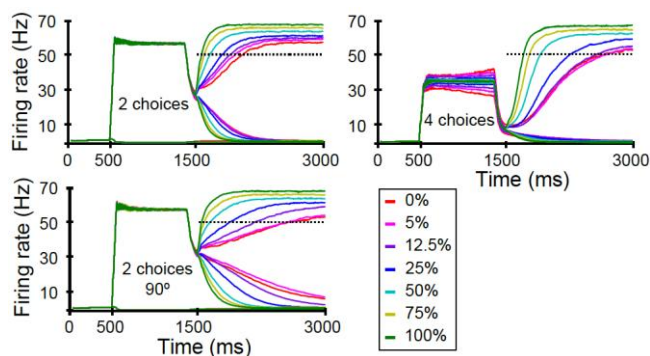
Fig. 4.4 Population averaged temporal evolution of firing rates at different motion strengths.
The average across all correct trials of each motion coherence level is shown for the three different experimental conditions. Different colors denote different motion coherence (see legend).

In sum, although all parameters in the network are independent of the number of possible alternatives, the attractor model can explain the lower firing rate prior to the activity build-up in the 4-choice condition due to the higher global inhibition. Longer reaction times in the 90°-case have a different origin: they are caused by stronger attraction of the symmetric state leading to a shallower slope in the 2-choice 90°-case compared to the standard 2-choice condition.

In the next section we will further elucidate the origin of the conditional differences in simulated behavior and neural activity with a mean-field analysis of the network's attractor states.

4.3.3 Mean-field approximation and range of decision-making

As already stressed, other than the respective number of pools receiving the target input, neither network parameters nor inputs in our model depend on the number of alternatives. The network thus exhibits categorical decision-making for two and four choices for the same range of external input v.

To investigate how this overlap of decision regimes for two and four alternatives depends on different network parameters, we used a mean-field approximation of the attractor model (Brunel and Wang, 2001) (see also 2.2.3 and A.1. With this approximation the computational cost of scanning the parameter space can be drastically decreased, as the number of dynamical variables is reduced to one for each neural population. By solving the mean-field equation one obtains the approximated average firing rate of each neural population when the system has settled into a stationary state. Consequently, starting from different initial firing rates, the fixed points of the firing rates can be calculated for the selective neural populations. There are three qualitatively different network states, whose existence and stability depend on the parameter configuration (Wong and Wang, 2006): the spontaneous state, with all pools firing at low rates; the decision state, where exactly one pool shows considerably higher activity than the others; and mixed states with two or more pools firing at high rates. The four different initial conditions we used cover the possible firing rates at different temporal stages of the spiking simulation. They are described in detail in Appendix A.1.5. The range of external inputs where, for all initial conditions and both experimental paradigms, a decision is reached, i.e. one pool lapses into an up-state of high firing rate, was termed "range of decision-making". It defines a region of multistable decision states.

To explore the effect of the relative size of the selective pools on the range of decision-making, we performed the mean-field analysis for different coding levels. The coding level is defined as the fraction of

excitatory neurons in the network selective for one target direction, and thus determines the relative size of the neural populations representing a specific choice (selective populations).

Fig. 4.5A shows an example of the stable fixed points, the attractors, of the firing rates of the selective populations for the parameters used in
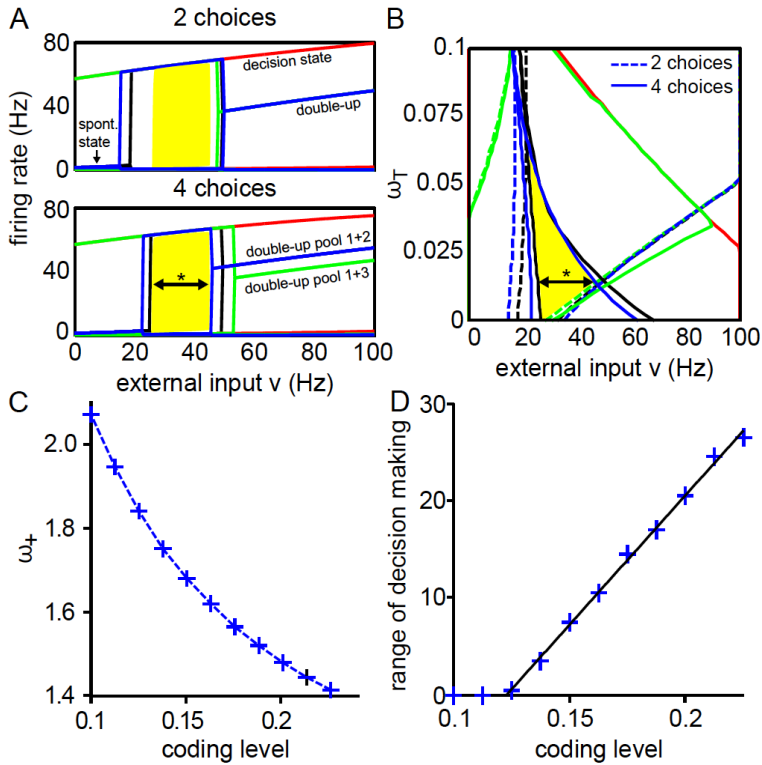
Fig. 4.5 Common range of decision-making for two and four alternatives in a mean-field approximation of the network.

(A, B) For the parameters used in the spiking model simulations (except $\omega_I = 1.1$ instead of 1.125), the stable fixed points were calculated with the mean-field approximation for the two- and four-choice condition and for four different initial conditions. Red: one pool starting at 120 Hz, the rest at 0 Hz. Green: two opposite pools 30 Hz, the rest 0 Hz. Blue: all four pools 30 Hz. Black: all four pools 0 Hz. (A) Stable firing rate fixed points ($\omega_T = 0.015$). (B) Starting and end points of decision states for the different initial conditions depend on the spatial connectivity ω_T between neighboring selective pools. (C) Dependence of the connectivity ω_+ on the coding level f in order to achieve equal fixed-point firing rates. (D) Optimal range of decision-making increases linearly with the coding level. For each coding level, the optimal ω_T was determined as shown in (B) and the broadness of the yellow region, the range of decision-making, was plotted against the coding level. The black line is a linear fit to the data.

the spiking simulations (coding level of 0.2, $\omega_+ = 1.48$ and $\omega_T = 0.015$). The different colors denote different initial conditions. At lower external inputs, for some initial conditions all pools stay in their spontaneous state, while for higher external inputs (from $\nu \approx 50$ Hz) a mixed "double-up" state emerges. In the decision state, exactly one pool is firing at a high rate, in the double-up state two pools. The traces of all other pools firing at low rates (≈ 2 Hz) overlap. For four alternatives all pools receive the input ν, which causes two possible double-up states: one with neighboring pools (e.g. 1+2) firing at high rates, the other with opposing pools (e.g. 1+3) firing at high rates. In the 2-choice case, the input ν is applied only to two opposing selective pools (1+3) and thus only these pools will fire at high rates.

Increasing the external inputs even more will result in mixed states with three and four pools firing at high rates, as observed for example during the target phase (not shown). The yellow regions in Fig. 4.5A and B depict the overlap of decision states for the two- and four-choice case, where for all initial conditions one pool wins the competition and a categorical decision is made (the range of decision-making).

Fig. 4.5B shows the dependence of the range of decision-making (i.e. the width of yellow region, coding level = 0.2) on ω_T, the additional spatial connectivity component between neurons from neighboring selective pools. As one can see in Fig. 4.5B, there is an optimal value of ω_T for a given parameter set (*black arrow). Fig. 4.5A is a horizontal cut through Fig. 4.5B at $\omega_T = 0.015$, the optimal value for a coding level of 0.2, as used in the spiking simulations.

When changing the coding level, the connectivity ω_+ has to be adapted as shown in Fig. 4.5C to keep the firing rates of the up-state at the values obtained for the spiking simulation parameters (≈ 60-80 Hz, Fig. 4.5A), which match the experimental observations (Churchland et al., 2008). The range of decision-making was then measured at its optimal ω_T value and plotted against the coding level (Fig. 4.5D). Interestingly, we found the range of decision-making to increase linearly with the coding level. For a coding level smaller than 0.125 no common decision state can be found for the 2- and the 4-choice task, regardless of the spatial connectivity ω_T. The linear relation of the coding level to the optimal range of decision-making was also found for a smaller AMPA/NMDA ratio than that used in the simulations, for the same and higher connectivity ω_+, confirming the generality of the outcome (Fig. 4.A.3).

4.4 Discussion

In this chapter, we presented a biophysically realistic spiking-neuron model for decision-making with two and four alternatives. Notably, all network parameters and inputs in our network are independent of the

number of possible alternatives we tested. Differences in firing rates and psychometric functions are solely based on the number of possible R-targets presented, which in the network corresponds to the number of pools receiving the target input. Moreover, we not only extended Wang's (2002) model to more than two choice-alternatives, but also analyzed how the size of the neural populations that encode the choice alternatives affects the network's capacity for multiple-choice decision-making. In a mean-field approximation of the network we found a linearly increasing relation between the relative size of the selective pools (the coding level) and the range of decision-making. This implies that pooling over many neurons favors decision-making independent of the number of choices.

4.4.1 Network properties and parameters

The presented network is an extension of Wang's (2002) model for binary decision-making to four alternatives. Like the original models (Brunel and Wang, 2001; Wang, 2002), it is capable of storing information by exhibiting persistent activity, because of slow recurrent excitation, which also enables the accumulation of sensory information (Fig. 4.A.2). Categorical decision-making in the network is based on attractor state dynamics and feedback inhibition, which mediates competition. Besides modifying the number of selective pools representing the possible choices from two to four, we introduced an additional spatial connectivity component ω_T between neighboring pools to model the circular location of the targets in the experiment, assuming a slightly higher correlation between pools 90º apart than between anticorrelated pools 180º apart. Quantitatively the connectivity between neighbors is only increased by 1.7% in our simulations, but ω_T proved essential to regulate the range of decision-making (Fig. 4.5B). It is generally advantageous to optimize the overlap between the competition regimes for different numbers of alternatives, as this minimizes the need for additional regulation mechanisms like top-down signals from other brain areas. An adaptation of the network connectivity during the learning phase to optimize the range of decision-making is therefore a plausible process.

4.4.2 Discrete or continuous representation?

Extending decision-making to more than two alternatives finally amounts to the question about continuous alternatives. At present, there are no experimental results on a RDM task with more than four discrete or continuous alternatives. Thus, it is still not clear when subjects reach their limit to distinguish possible motion directions, or how accurately they can determine motion direction in a continuous task. Infinite precision may

not be needed to obtain the final resolution of the cognitive and motor systems (Churchland et al., 2008). A discrete network model with a finite number of selective neural populations might thus be sufficient to account even for continuous choices.

In the following, we will discuss our results in comparison with the recently proposed continuous models (Beck et al., 2008; Furman and Wang, 2008) of multiple-choice decision-making. Furman and Wang's (2008) model, like ours, is a biophysically detailed attractor model. Their network consists of directionally tuned neurons whose preferred directions cover a full circle. Excitatory neurons are connected according to a Gaussian curve depending on their difference in preferred directions. This continuous approach allows for testing eight R-targets spaced 45° apart. However, for eight alternatives, in 49% of the trials no categorical decisions could be made, because activity around adjacent R-targets tended to merge. In addition, their continuous model could not account for the differences in reaction time observed between the standard 2-choice- and the 90°-case (Churchland et al., 2008). The two conditions resulted in identical outcomes in their simulations (Furman and Wang, 2008).

Our discrete model with spatially tuned connections between the selective pools could account very well for all tested paradigms, including the 90°-case. Additionally, we were able to explain the intermediate RTs of the 90°-case, based on a prolonged symmetric state of the two selective populations. Thus, pooling over many neurons and introducing a graded spatial connectivity between the pools might represent the physiological conditions of neurons in LIP more accurately than a ring structure where each neuron encodes one particular direction.

Beck et al. (2008) took a distinct, parallel approach with respect to the biophysically realistic attractor models. Their model is focused on possible probabilistic properties of neurons. It fits Churchland et al.'s (2008) data well. However, in contrast to the experimental findings, different activity thresholds were used in the 2- and 4-alternative case to terminate the decision. The probabilistic approach so far accounts well for single cell data. Yet, to verify if populations of LIP neurons really encode probability distributions as predicted, future multiunit recording experiments are required (Beck et al., 2008).

4.4.3 The importance of pool size

Apart from successfully accounting for all experimental paradigms, the pool structure of our model extension entails an important advantage regarding the network's dynamics: the presented network operates in a range of categorical decision-making, independent of the three tested conditions. In Furman and Wang's (2008) model, the simulations needed to be controlled by inputs dependent on the number of possible R-targets,

which were supposed to originate from unknown higher level brain areas or normalization processes. They introduced an adaptation of the target input and an external control signal during the decision-making period as mechanisms to regulate the dynamical range of decision-making for the respective numbers of alternatives.

In contrast, in the discrete model there is no need to adapt neither parameters nor inputs to the number of choice-alternatives. Using a mean-field approximation, we found that the overlap of decision regions for two and four choices increases linearly with the relative size of the selective populations that encode the choice-alternatives (the coding level). A possible explanation of this effect might be given based on the recurrent connectivity. To keep the fixed-point firing rates of the model constant while increasing the coding level, the recurrent connectivity of the network has to be adapted nonlinearly, as shown in Fig. 4.5C. In turn, if more neurons encode a choice-alternative, each neuron in that pool will receive higher recurrent inputs. In fact, we observed that the recurrent activity a single neuron receives increases linearly with the coding level. We believe that this increase of recurrent activity stabilizes the decision state.

Furman and Wang (2008) suggested that their top-down signal could also serve to control the speed-accuracy tradeoff. In the discrete network an input-based control of the speed-accuracy tradeoff could be implemented with the advantage of being independent of the number of possible choices, following the study of Roxin and Ledberg (2008), where increasing the inputs was found to decrease reaction times and performance monotonically.

Although the network structure of Furman and Wang's (2008) model generally offers the possibility to simulate any number of alternatives and angular locations, a biophysically realistic model accounting for all experimental paradigms and more than four choice alternatives is still an objective for future research.

Taken together, our results indicate the benefit of a pooled, multi-neuron representation of choice-alternatives. We suggest that a bigger pool size, apart from the obvious advantage of redundancy in case of neural loss, or of averaging out noise by pooling across inputs, enables the network to exhibit relatively stronger recurrent inputs useful for stabilizing decision states.

4.A Chapter appendix

4.A.1 Network simulations without target stimulus

To investigate the network's behavior without the target input, we performed two additional sets of simulations: one with a motion signal of 20 Hz and one with 45 Hz applied to the four selective pools. All other network parameters are identical to the main simulations in the results section (4.3) of this chapter.

For a 20 Hz motion signal (Fig. 4.A.1B) and small coherence levels (up to 12.5%), the input to the selective populations is too low to induce decisions. In most of the trials with small coherences, all pools stay in their spontaneous states (see also Fig. 4.5A, main text: The black line represents the mean-field initial condition where all pools start at 0 Hz. For 20-Hz external input, the network stays in its spontaneous state). For
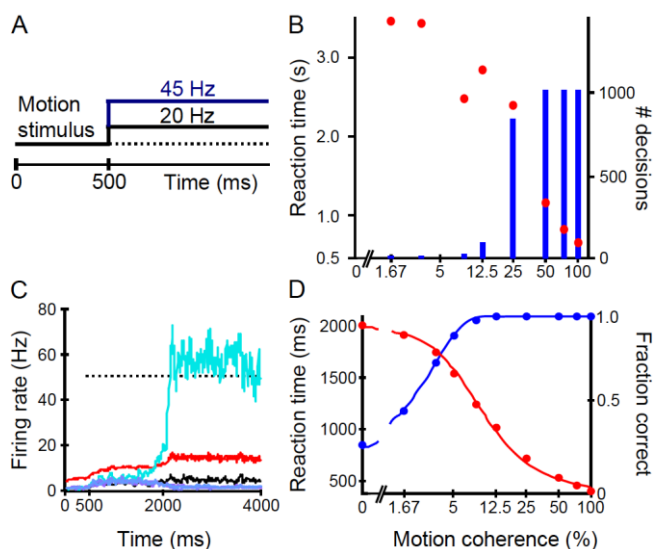
Fig. 4.A.1 Network simulations without target stimulus.
(A) Time course of external input. (B) 20 Hz motion input. The number of trials with decisions is shown with the blue bars (right ordinate). Reaction times are displayed as red dots (left ordinate). (C) Single trial at 0% coherence for a motion signal of 45 Hz. The integration process starts at 500 ms, a decision is reached at ~2,000 ms. (D) Speed and accuracy of the 45-Hz simulations. Red (left ordinate): reaction times of correct trials as a function of motion coherence, fitted by a hyperbolic tangent function (Eq. A.35). Blue (right ordinate): Psychometric function, fitted by a Weibull function (Eq. A.36). 1,000 trials were simulated for each data point.

higher coherence levels, however, the model is capable of reaching a decision in almost all trials and with realistic reaction times. At these high coherence levels, the model only takes correct decisions.

Fig. 4.A.1D shows the simulated behavior with 45 Hz motion input. For lower coherences, decisions take longer than with the target input (compare Fig. 4.2). This is realistic, because the random-dot motion task is more difficult without knowledge about the possible R-targets. Accuracy starts at chance level, as with the target input, and reaches 100% for higher coherence levels. At low motion coherence, at most 4% of the trials had to be excluded because they failed to reach a decision within the simulation time of 4,000 ms.

Our simulations show that for a strong enough motion signal, our network is capable of decision-making without the target input. This corresponds to possible experiments where subjects have to decide on the direction of coherent motion without being shown any R-target possibilities. However, it must be noted that our model is restricted to its four selective populations. To model continuous decision-making, a higher number of pools would be needed (see 4.4.2).

4.A.2 Attractor network is capable of persistent activity

Due to the strength of its recurrent connections, the presented attractor model is generally capable of displaying persistent activity. For the parameter set used to fit the experimental data of Churchland et al. (2008), the selective population that ''won'' the competition is able to sustain its activity at about 48 Hz in all trials, even if all selective inputs are switched off at 4,000 ms (Fig. 4.A.2).

In Fig. 4.A.2C the stable attractor states of the network calculated in the mean-field approximation are shown for different values of the recurrent connectivity ω_+. Persistent activity can be sustained by the network, if the decision state is stable at 0 Hz, that is without any selective external inputs to the selective populations. From $\omega_+ = 1.44$, the decision state exists and is stable even for zero external input. For smaller values of ω_+, a decision state with high firing rates occurs only at higher external inputs. In the main simulations, we used $\omega_= = 1.48$ and, therefore, reliably obtained persistent activity. For the mean-field calculation shown in Fig. 4.A.2C one pool was initialized at 120 Hz, the other three pools at 0 Hz, which corresponds to the situation when one pool has already ''won'' the decision.
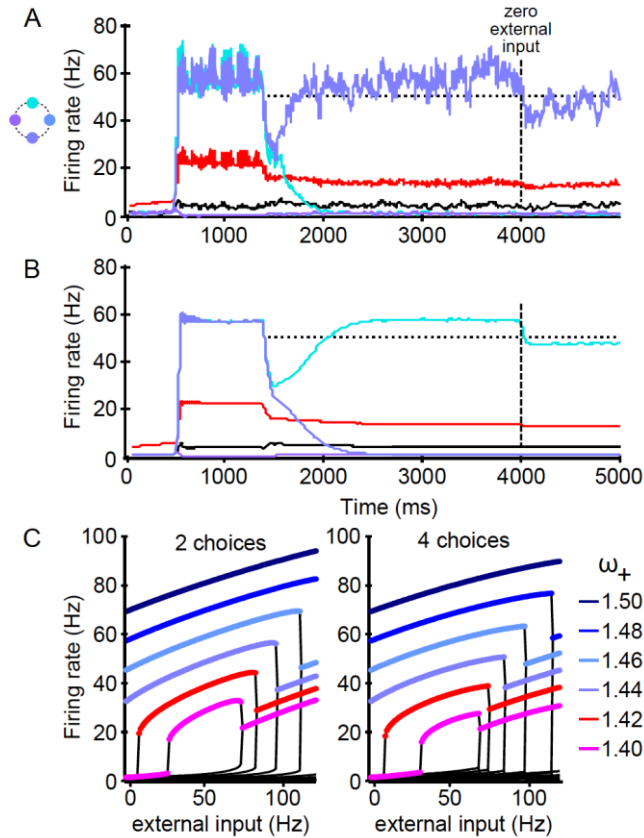
Fig. 4.A.2 Network exhibits persistent activity after the external inputs are switched off.
(A and B) Displayed are a single trial (A) and the trial average over 1,000 simulated trials (B) at zero motion coherence for the 2-choice case (colors as in Fig. 4.3; red and black lines denote inhibitory and nonselective pools). (C) Stable attractor states of the network were calculated in the mean-field approximation (A.1.5b) for different values of the recurrent connectivity ω_+.

4.A.3 Range of decision-making for smaller AMPA/NMDA ratio.

A linear relation between the range of decision-making and the coding level was also found for the original AMPA/NMDA ratio used in the model of Brunel and Wang (2001) ($\delta = 0$ in Eq. A.10). We tested two conditions: first, all network parameters were identical to those used in Fig. 4.5 and the spiking simulations, except for a smaller AMPA/NMDA

ratio (Fig. 4.A.3A,B). In this case, the up-state fixed-point firing rates are lower than the firing rates found in the experiments. The connectivity ω_+ was adapted as in Fig. 4.5C, so that the up-states deviated by less than 2 Hz from the up-state firing rate value at $\omega_+ = 1.48$, the connectivity weight used in the spiking simulations. The optimal range of decision-making still increases linearly with the coding level (compare Fig. 4.5D).

Next, we again used the original AMPA/NMDA ratio, but increased the connectivity ω_+ in order to increase the up-state firing rates to about the same range as in the spiking simulations and experiments (Fig. 4.A.3C,D). The range of decision-making again increased linearly with the coding level. However, due to the higher ω_+ values, the resulting optimal range of decision-making was much smaller, and already for a coding level of 0.175, no common decision state could be found for the 2- and 4-choice task.
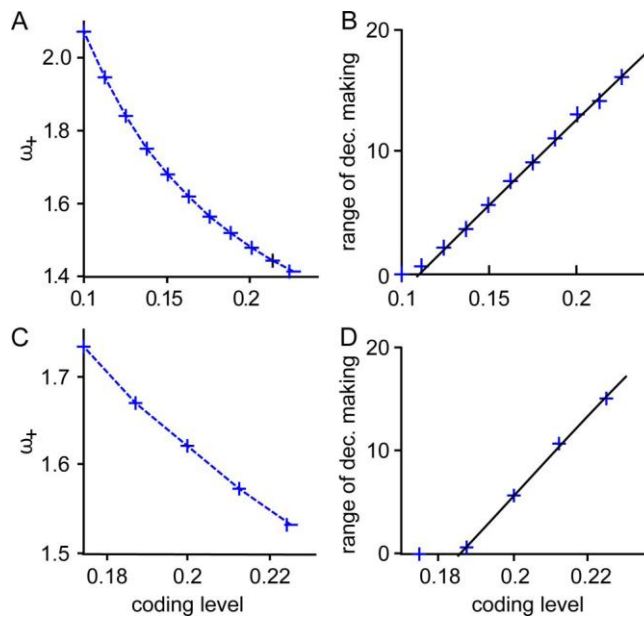
Fig. 4.A.3 Range of decision-making for smaller AMPA/NMDA ratio.
(A,B) Smaller AMPA/NMDA ratio. The up-state firing-rates were adapted to the value at $\omega_+= 1.48$. (C,D) To compensate the smaller AMPA/NMDA ratio, the recurrent connectivity was increased, in order to fit the experimentally observed firing rates with the smaller AMPA/NMDA ratio. Increasing ω_+ leads to a much smaller range of decision-making.

5 A MULTIPLE-CHOICE TASK WITH CHANGES OF MIND

The work presented in this chapter is currently prepared for submission. An abstract was published at the SfN 2010 *conference[13].*

5.1 Introduction

Compared to the decisions we are faced with every day, decision-making in psychophysical experiments is typically reduced to the highly simplified conditions of two-alternative forced-choice (2AFC) tasks (Luce, 1991; Bogacz et al., 2006). Nevertheless, the behavioral data gained from these perceptual tasks, together with complementary evidence from single-cell recordings (Roitman and Shadlen, 2002; Ratcliff et al., 2003), motivated and constrained formal models of decision-making (Smith and Ratcliff, 2004; Wang, 2008).

Still, 2AFC tasks, by definition, neglect important features of real-life decision-making such as choices between multiple alternatives and choice reevaluation after an initial decision. As we have seen in the last two chapters, several authors have recently extended the RDM paradigm and investigated these more complex aspects of decision-making independently. On one hand, Churchland et al. (2008) and Niwa and Ditterich (2008) augmented the RDM task from binary to multiple choices. In particular, Churchland et al. (2008) tested monkeys on a 4-alternative RDM discrimination task and compared their behavioral and neurophysiological responses to the standard binary task. On the other hand, Resulaj et al. (2009) considered "changes of mind" which are thought to arise through further processing of available information after an initial decision has been made. Instead of a saccadic response to the chosen R-target, as in the standard RDM task, human participants had to move a handle to a left or right response target. The continuous hand movement allowed observing the subjects' behavior after their first decision. Indeed, participants occasionally switched from one direction to the other, thus "changing their mind".

Here, we combined these two paradigms in order to explore how changes of mind depend on the number of choice-alternatives (Fig. 5.1). When presented with four possible alternatives, our naive human participants committed more initial errors (i.e. choosing the wrong

---

[13] Albantakis L, Branzi FM, Martin C, Costa A, Deco G (2010) Changes of mind during binary and multiple-choice decision-making. 40th Annual Meeting Society for Neuroscience: SfN*2010 San Diego. Poster 503.5/KKK55

response) than in 2-choice trials. A higher number of possible motion directions leads to more confusion between choices. Therefore, one might also anticipate more changes of mind for four than for two alternatives. Yet, the percentage of trials with changes of mind did not significantly differ for the different numbers of alternatives. In fact, slightly fewer changes were observed in the case of four possible choices. Particularly the probability of "correcting" changes of mind, from the wrong to the correct choice, was lower for four than for two alternatives.

Furthermore, our experimental results could be fit by an attractor model of decision-making. This model was constructed by merging the two biophysically realistic attractor models for decision-making presented in the previous two chapters, which explained choice behavior, neural activity and changes of mind of the two preceding experimental studies (Albantakis and Deco, 2009, 2011). What is more, the attractor model could also account for between-subject variations by an alteration of the decision threshold. A lower threshold in the model caused faster reaction times and more changes of mind. This corresponds to the negative correlation we found between the mean reaction times of individual participants and their tendency to change their minds. Finally, the computational model suggests, that extensive training might enhance behavioral differences in reaction time and changes of mind between the different experimental conditions. This prediction could be tested with over-trained human participants and maybe also with monkeys trained on the 2- and 4-choice task.

Our main results are:

- Choice corrections are less probable for more choice alternatives.

- Across individuals changes of mind correlate negatively with mean reaction times.

- An attractor model can explain the changes of mind dependence on condition and reaction time.

- We predict that extensive training might enhance differences between conditions.
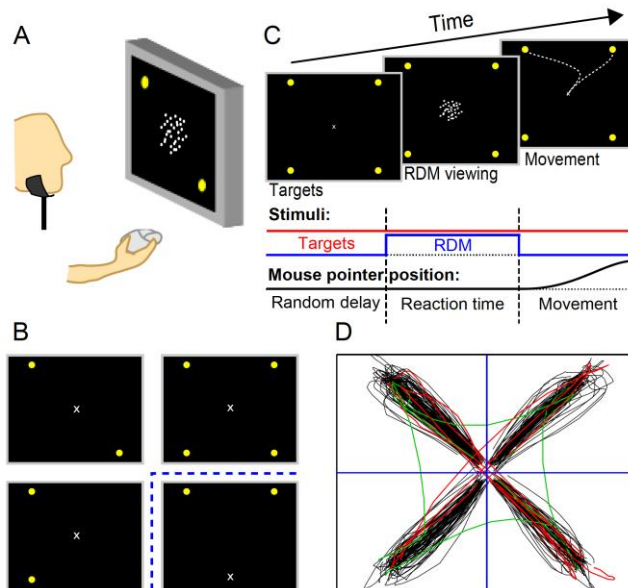
Fig. 5.1 Experimental paradigm: setup, time course and conditions.
(A) Participants had to indicate their choice on the net direction of motion by moving a computer mouse pointer towards the respective visual R-target on the computer screen in front of them. Yellow dots denote the R-targets. (B) Illustration of the three possible R-target arrangements in the 2- and 4-choice experiment and the 2-top control. (C) Experimental time course. (D) Example traces from one participant (4-choice condition). In the majority of trials the subjects moved directly to one of the visual R-targets (black traces). Some trajectories, however, revealed a change of mind during the movement: they started towards one, but terminated at another R-target. Changes could be observed between adjacent (green) and opposite (red) R-targets.

5.2 Methods

5.2.1 Experimental paradigm

In the following we will describe the general experimental procedure of the 2- and 4-choice RDM task with changes of mind. Further details can be found in Appendix A.3.

a) Experimental task and visual stimuli

The random-dot motion discrimination task is illustrated in Fig. 5.1. In this task, 15 human participants were presented with a RDM stimulus.

The stimulus varied in the percentage of direction-coherent moving dots. Participants were asked to decide as fast and accurately as possible on the net direction of motion in the stimulus. They had to indicate their choice by moving a mouse pointer to the related response target (R-target) aligned with the identified motion direction. In each trial, there were either two or four possibilities for the direction of coherent motion.

Each trial began with the presentation of a gray circle of 2º diameter in the screen center with a small fixation-cross in the middle (the so-called "start-target"). To start the trial, participants had to click on the start-target with the mouse, and the potential response options (R-targets) appeared on the screen. The yellow R-targets indicated the possible directions of coherent motion. They could appear either:

a) in each of the four corners of a virtual square (4-choice trials), or

b) in just two of the four corners (2-choice trials).

The two presented R-targets could be either 180º apart, and thus symmetrically located with respect to the fixation-mark, or 90º apart, and therefore contiguous on one side of the screen (Fig. 1B, top and bottom left panel respectively). The R-targets remained present on the screen until the end of the trial.

Due to the R-targets, participants knew prior to the onset of the motion stimulus how many choice-alternatives they had and which coherence directions were eligible for each given trial. After a random delay[14], the fixation disappeared and was replaced by the RDM display. Upon the presentation of the RDM stimulus, participants were free to respond by moving the mouse towards the preponderant direction of motion, and clicking into the corresponding R-target. Until they decided to respond, participants had to keep the mouse pointer within a 2º diameter around the fixation-cross. The RDM stimulus was extinguished as soon as subjects moved the mouse device out of this area. The time limit to leave the starting position was 2 s (time-out #1). Once they moved out of the starting position, the subjects had 1 s (time-out #2) to click on one of the R-targets[15]. Only trials where one of the R-targets had been clicked within these time-outs were counted as "valid trials".

b) Experimental sessions

Our participants underwent four experimental sessions of 30 minutes each, all the same day, separated by a time interval of two hours. In the first three sessions, we tested the participants on the combined 2- and 4-alternative task explained above. In these sessions trials with two and four alternatives were randomly presented, but with the same total number of

---

[14] The delay time was sampled from a truncated exponential distribution (range 0.7–1.0 s; mean 0.82 s).

[15] More precisely they had to hit a square area of 3º edge length around the R-target.

trials for each of the three conditions (2-choice 180º, 2-choice 90º, and 4-choice condition). We used a set of eight different coherence levels (0%, 3.2%, 6.4%, 12.8%, 25.6%, 51.2%, 76.8% and 100%), which were presented 16 times each, except for 0% which was presented four times[16]. This resulted in a total of 348 trials per session. After the first half of trials participants could have a small break. In the beginning of each session we also included 87 practice trials to familiarize the participants with the task and to assure a satisfactory level of performance.

In the fourth experimental session, we replicated the visual arrangement used by Resulaj et al. (2009) as a control to compare our simpler setup with the original binary changes of mind paradigm ("2-Top control"). Here, participants had to decide whether the coherent motion was horizontally right or left. Hence, in this fourth session participants were always presented with the same two possible motion directions and the following R-target configuration: the start-target was located at the bottom, 14º below the center, and the two R-targets at the top of the screen, 19.8º from the center, on the edges of a virtual triangle (Fig. 5.1.B, bottom right panel). Accordingly, in this configuration, the R-targets were not perfectly aligned to the possible motion directions. In the 2-top control we used only the first six coherence levels, and a total of 440 trials preceded by 88 practice trials, time-out #1 was 2 s, time-out #2 was 0.7 s. Except for that, the 2-top control was conducted in the same way as the 2- and 4-choice task described above.

Task instructions were always first given verbally and afterwards repeated visually on the video-screen. The subjects were not explicitly informed that they could change their decision during the motion response.

c) Data analysis

In the analysis of our experimental data, three dependent variables were of special interest. First, we measured reaction times (RT), corresponding to the time taken by the participants to initiate their response. More precisely, it denotes the time between the onset of the RDM stimulus and the moment when the mouse pointer left the start-target area. Second, we measured the accuracy of the responses, i.e. the percentage of correct trials. The percentage of trials in which the correct R-target was selected determined the "final performance" of the participant. Third, and more importantly for our purposes, we measured the mouse pointer trajectory at 75 Hz. Based on the mouse pointer trajectories, we determined whether subjects changed their initial decision. If a trajectory was initially directed towards one R-target, but then crossed

---

[16] To avoid frustrating participants with unsolvable trials, the 0% coherence level was only presented four times for each R-target condition. For 0% coherence the "correct" target was defined randomly.

the horizontal or vertical axes and ended at another target, this trial was counted as a change of mind. As in Chapter 3, we further distinguished between "correcting changes", which were responses that started towards an incorrect R-target, but then turned to the correct R-target, and "erroneous changes", which turned from the correct to a wrong choice. The initial direction of the trajectory was interpreted as the participants' "initial performance". Trajectories of trials without changes were typically straight to the target (see Fig. 5.1D). The initial direction and changes of mind could thus be determined reliably by the deviation from the diagonal between the center and the chosen R-target (see Appendix A.3.2).

In this chapter, we focus on comparing the performance and changes of mind between two and four alternatives. Therefore, in the analysis of our experimental results, we collapsed all 2-choice trials across the different target locations, and the 90°- and 180 °-cases. Initially, we included the 90°-case as a control condition to tell apart whether differences in behavior indeed resulted from the different number of alternatives, or rather from the smaller angular distance in the 4-choice case (Churchland et al., 2008). Nevertheless, we combined both 2-alternative conditions here, because in the 4-choice condition confusion between different alternatives could also occur between adjacent and opposite R-targets. For a fair comparison, the 2- choice trials should thus also include both angular distances (90° and 180°). The minor differences we found between the 2-choice 90°- and 180°-cases will be treated in Chapter 6.

5.2.2 Computational model

In Chapter 4 we showed that biologically-inspired attractor models can account for primate decision behavior and neural activity in a multiple-choice RDM task. Furthermore, they also match human choice behavior and changes of mind in a binary RDM paradigm (Chapter 3). In this chapter, we assess the conformity of a modified version of the multiple-choice model presented in Chapter 4 with the present set of psychophysical results. The network characteristics and kinetics are again summarized in Appendix A.1 and Table A.5. All default simulation parameters are listed in Table A.6.

a) Network structure and connectivity

In short, the network consists of 500 leaky integrate-and-fire neurons with conductance-based synaptic responses (400 excitatory pyramidal cells and 100 inhibitory interneurons).

The network structure is sketched in Fig. 5.2A. Excitatory neurons are subdivided into four selective populations or "pools" (each 80 Neurons), encoding the four possible motion directions, and a fifth pool of
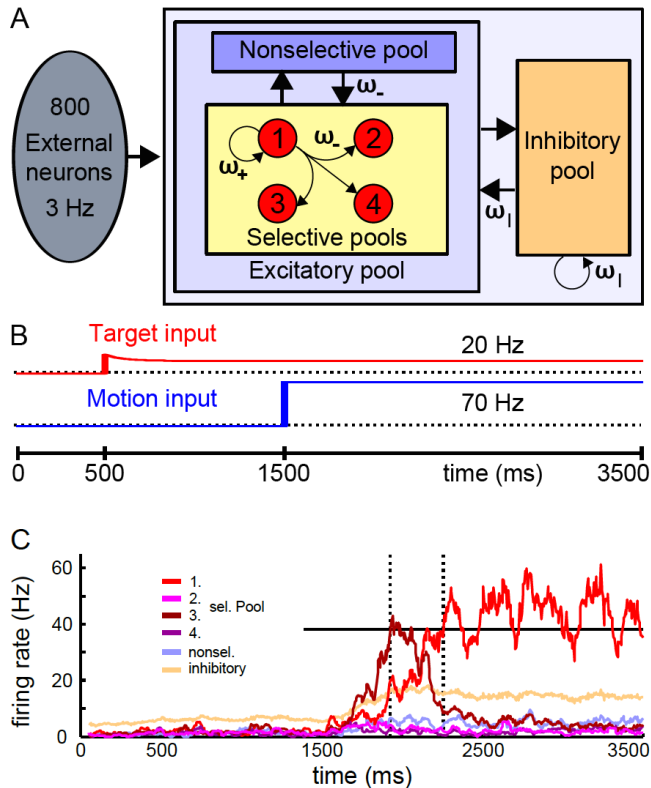
Fig. 5.2 Computational model: populations, connectivity and input.
(A) Diagram of the attractor model for decision-making between up to four choice alternatives. Unlabeled arrows denote a connectivity of 1 (baseline). (B) Time course of target and motion input to the selective populations in order to model the experimental design of the RDM task. (C) Example trial with "change of mind" for two alternatives at 12.8% coherence. The initially winning population (first threshold crossing) is overtaken by the other firing rate transient. The horizontal black line at 38 Hz indicates the threshold. Dotted vertical lines mark times of threshold crossings.

nonselective neurons. The latter emulates the activity of surrounding neurons that are not selective to any of the four R-target directions. All excitatory neurons are connected to one pool of inhibitory neurons, which regulates the overall activity by implementing competition in the network.

The recurrent connections between neurons within one selective pool are stronger ($\omega_+ = 1.48$) than between selective pools and from the nonselective to selective pools ($\omega_- = 0.88$). Other than in the multiple-choice model presented in Chapter 4, there is no spatial connectivity component here (discussed in Section 5.5.3).

As before, we used a mean-field reduction of the full spiking-neuron model to initially locate the working point of the network with respect to the two crucial bifurcations that contain the range of categorical decision-making (Fig. 2.8 and Fig. 3.6).

b) Simulation of sensory inputs

Three types of external inputs were applied to the neural network as noisy uncorrelated Poisson spike trains. First, all neurons received a background input of $\nu_{ext} = 2.4$ kHz, equivalent to 800 excitatory connections from external neurons firing at 3 Hz. Furthermore, two "sensory" inputs mimicking the task-relevant visual stimuli, namely the R-targets (ν_{target}) and the random-dot motion stimulus (ν_{motion}) were applied to the selective populations only (Fig. 5.2B).

Target input. As in Chapter 4, the target input ν_{target} was applied to two of the selective pools in the 2-choice condition, and to all four selective pools in the 4-choice condition, corresponding to the experimental R-target stimuli. Departing from the target input used in the last two chapters, here we did not include the initial phase of high inputs. Instead, the target input was mostly constant in time, with only modest initial adaptation:

$$\nu_{\text{target}} = 20 \text{ Hz} + 10 \text{ Hz} \cdot \exp(-(t - 500 \text{ ms})/\tau), \quad (5.1)$$

where 500 ms is the onset time in the simulation and $\tau = 100$ ms the adaptation decay time constant. We will comment on this and other differences between the present and previous implementations of the attractor model in more detail in the results and discussion sections.

Motion input. Other than the target input, the RDM input ν_{motion} was always applied to all four selective populations for the 2- and 4-choice condition. Coherent motion was simulated as a positive bias to one selective pool, balanced by a reduction of the motion input in the other three selective pools:

$$\nu_{\text{motion,1}} = 40 \text{ Hz} + 30 \text{ Hz} \left(1 + \frac{3 \cdot coh}{100\%}\right), \quad \nu_{\text{motion,2-4}} = 40 \text{ Hz} + 30 \text{ Hz} \left(1 - \frac{coh}{100\%}\right). \quad (5.2)$$

Thereby, the total motion input to the network was kept constant. A motion coherence of 100% thus corresponds to a bias of 90 Hz to one selective pool, resulting in $\nu_{motion,1} = 160$ Hz applied to the first selective pool and $\nu_{motion,2-4} = 40$ Hz to the other selective pools.

c) Decision threshold and network simulations

A (first) decision was reached in the simulations, when the activity of one selective pool crossed a fixed decision threshold and surpassed the other pools by at least 5 Hz. The same conditions applied for a change of mind. The threshold was assumed to be independent of motion coherence and the number of choice alternatives (Roitman and Shadlen, 2002;

Churchland et al., 2008). To fit the mean behavioral data averaged across all 15 participants, we used a threshold value of 38 Hz. To account for the participants' behavior grouped according to their total number of changes, we respectively used 32 Hz, 38 Hz and 42 Hz for the groups with most, middle, and least changes on the same simulated trials. All threshold values were selected by hand with a resolution of 1 Hz, based on the overall best fit to experimental reaction times, performance, and changes of mind.

Reaction times were calculated as the time of threshold crossing plus a non-decision time $t_{ND} = 350$ ms (Niwa and Ditterich, 2008), which also set the time limit for changes of mind (see Section 3.2.2 for a detailed discussion of this time-out). A t_{ND} of 350 ms is in accordance with an assumed afferent signal latency of about 200 ms for the motion signal (Roitman and Shadlen, 2002; Churchland et al., 2008), plus 150 ms to account for movement initiation and execution (Cui and Andersen, 2007; Snyder et al., 1997).

5.3 Experimental results

To study the relation of changes of mind to the number of choice alternatives, we asked naive human subjects to perform a random-dot motion (RDM) discrimination task with two and four possible directions of motion (Fig. 5.1).

In addition, we intended to replicate preceding findings on changes of mind (Resulaj et al., 2009) with our simpler experimental setup. Therefore, we ran a separate block of trials with two alternatives in the "2-top" control condition, where the response targets (R-targets) were arranged as in (Resulaj et al., 2009). The results of the 2-top control are summarized in Fig. 5.A.1 and will be compared to the main 2- and 4-choice experiment, in the discussion.

In all experimental conditions, we recorded the participants' choices, the time between the onset of the motion stimulus and movement initiation (reaction time or "RT"), and the movement trajectories of the mouse pointer. Changes of mind were generally defined as those traces starting towards one R-target but then changing the direction and crossing the horizontal or vertical axes. Such trajectories could be observed occasionally in all experimental conditions. Notably, in the 4-choice condition subjects not only changed between adjacent R-targets, but also across the diagonal (Fig. 5.1D). In the following we will first report the pooled data from all 15 participants, but will show subgroups and individual results later.
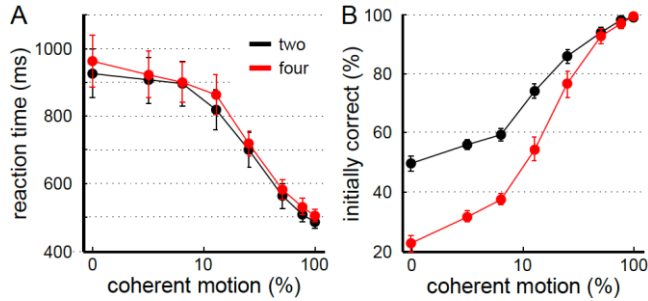
Fig. 5.3 Mean reaction times and initial performance.
(A) Reaction times decreased with increasing coherent motion and were longer for four alternatives. (B) Initial performance, that is the percentage of initially correct choices, started at chance level (50% for two and 25% for four choice alternatives) and increased to almost perfect accuracy for 100% coherent motion.

5.3.1 Reaction times and choice accuracy

As shown in Fig. 5.3, the RTs in correct trials decreased as motion coherence increased. This effect ranged from about 950 ms for the most difficult coherence condition (0%) to 500 ms for the easiest one (100%) (Fig. 5.3A). In other words, the more dots move in the same direction the faster the response is taken. More interesting perhaps is the fact that participants needed more time to decide in the 4-choice than in the 2-choice trials.

To test this statistically, we performed a within-subject ANOVA on the reaction times of all correct trials (collapsing the three sessions), excluding changes of mind, with the factors "number of choices" (two vs. four choices) and "coherence" (eight levels of coherence: 0%, 3.2%, 6.4%, 12.8%, 25.6%, 51.2%, 76.8% and 100%). The results revealed a main effect for both, "coherence" ($F_{(7,98)}=45.88$, $p<.001$) and "number of choices" ($F_{(1,14)}=22.03$, $p<.001$). Moreover, the interaction between "number of choices x coherence" was also found to be significant ($F_{(7,98)}=2.39$, $p<.05$). The difference between 2- and 4-choice trials was specifically significant for the following coherence levels: 0%, 12.8%, 76.8% and 100% ($p<.05$). Besides, the difference between RTs in the condition without coherent motion (0%) and other coherence levels started to become significant with 25.6% coherent motion ($p<0.5$).

Coherence also affected accuracy in a similar way: the higher the coherence, the higher the percentage of initially correct choices ("initial performance") (Fig. 5.3B). As expected, accuracy started at chance level for 0% coherent motion and reached close to perfect performance for the

highest coherence levels. Also at intermediate coherence levels accuracy was lower in the 4-choice than in the 2-choice condition, which reflects the lower prior probability of each choice given four alternatives.

5.3.2 Changes of mind

As explained above, changes of mind were defined as mouse trajectories that were initially directed towards one R-target, but then turned to another. They occurred in all experimental conditions, at all coherence levels, and could lead to a correct or an incorrect final choice. In Fig. 5.4 A and B, correct and erroneous changes of mind are plotted for each motion coherence level as percentage of all valid trials. Intuitively, changes to the correct response (Fig. 5.4A) were most frequent at intermediate difficulty (25.6%). This is because for low coherences, there is only little sensory evidence which might induce a change from the initial decision. Then again, at high coherences participants already initially chose the right direction of motion. Hence, only intermediate conditions left room to find a substantial number of changes (Resulaj et al., 2009). Erroneous changes generally decreased with increasing motion coherence[17] (Fig. 5.4B).

Considering all coherence levels, participants changed somewhat more in the 2-choice than in the 4-choice condition (black and red lines in Fig. 5.4A,B respectively). This difference is more evident if we consider it in the context of the different accuracy levels of the two conditions. Given four possible motion directions and low coherence levels, our participants committed about twice as many initial errors than in the 2-choice condition (Fig. 5.3). They might thus also have corrected their choice more often. Yet, we observed the opposite. Fig. 5.4C and D show the same data as Fig. 5.4A and B, but here we considered accuracy (i.e. the fraction of correct choices). Instead of dividing by the total number of valid trials, the number of changes was respectively divided by the initial errors or initially correct trials. More precisely, the correcting changes were divided by the number of initial errors from which they could have originated. Similarly, the erroneous changes were divided by the number of initially correct trials. Plotted in this way, it becomes clearer that the probability to make a correcting change was actually much lower when subjects were presented with four alternatives than for just two. However, the probability of changing from the initially right choice to a wrong R-target (Fig. 5.4D) was very similar in the two conditions.

---

[17] Note that at 0% coherence the total number of changes was divided by the number of choice alternatives. As without coherent motion there was no initially right or wrong choice, changes there could neither be counted as correcting nor erroneous.
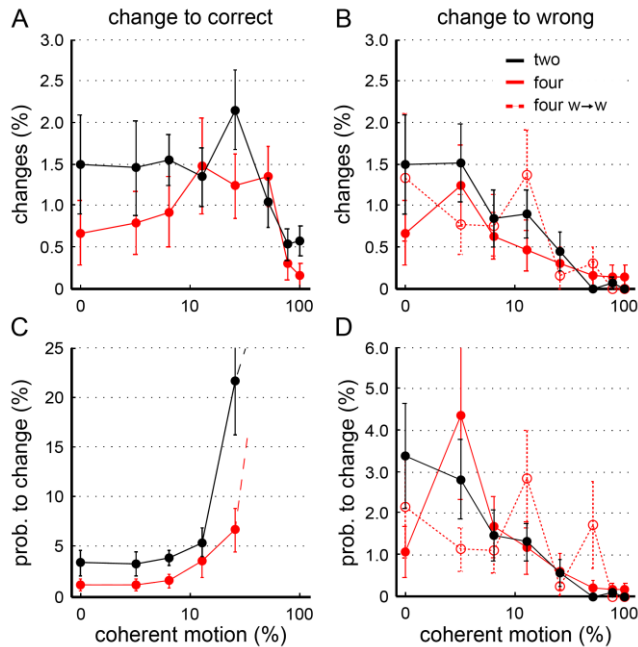
Fig. 5.4 Comparison between changes of mind for two and four alternatives.
(A) Correcting changes and (B) erroneous changes are shown as percentage of all valid trials. Given four choice-alternatives, also changes between two wrong R-targets occurred (w → w, dashed line in B). (C,D) The same data as in (A,B), but considering the different accuracies (see Fig. 3B). Correcting changes (C) and changes from wrong to wrong (dashed line B,D) were divided by the number of initial errors, changes from correct to wrong by the number of initially correct trials. Thus (C) and (D) show the probability to change from an initially correct or from an initial erroneous choice respectively. Note that we plotted changes scaled by initial errors (C) only up to 25.6 % coherence, because for higher coherent motion the number of initial errors became very small (leading to division by small values and thus uninformative probabilities).

Finally, in the 4-choice condition, participants also changed between two wrong alternatives (Fig. 5.4B,D dashed line). Except for noise fluctuations the visual stimuli neither presented evidence for, nor against these changes. Note also that wrong-to-wrong changes do not affect the final performance.

All in all, changes of mind nevertheless improved participants' accuracy (Fig. 5.5), with the exception of the lowest coherent motion level (3.2%) in the 4-choice condition. There, the higher number of erroneous changes compared to correcting changes could be expected. That is because there are more distracters than correct R-targets and the evidence
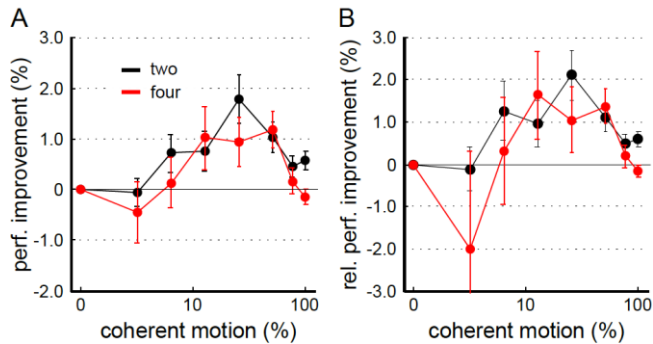
Fig. 5.5 Performance improvement through changes of mind.
(A) Absolute difference of initial and final performance (considering changes of mind). (B) Relative performance difference, i.e. the performance difference shown in (A) divided by the initial performance.

for changing was weak compared to the noise fluctuations in the visual stimuli.

In summary, there are two empirically relevant observations: first, erroneous changes of mind are equally probable in the 2- and 4-choice condition. The probability to dismiss a correct choice decreases with higher coherence, but is apparently not influenced by more distracters. Second, correcting changes of mind do depend on the number of possible choices. They are less frequent for more alternatives.

5.3.3 Correlations between changes of mind and mean RT for individual participants

Another interesting observation in our study refers to the individual differences in the tendency to make changes of mind. These individual differences can only be detected in experiments in which a relatively large sample of participants is tested and, perhaps, for that reason have not been addressed in detail in previous studies (Palmer et al., 2005; Niwa and Ditterich, 2008; Resulaj et al., 2009).

Participants who made more changes tended to have faster reaction times. This can be appreciated in Fig. 5.6A, where the overall number of changes (ONC) is negatively correlated with the mean reaction time of all valid trials (Pearson correlation coefficient $R = -0.63$, $p < 0.05$, N=14, excluding the one outlier).

Furthermore, there was a trend towards a positive correlation between overall accuracy and reaction times (Fig. 5.6B), namely the slower the responses the higher the accuracy ($R = 0.48$, $p = 0.085$, $N = 14$). This trend is consistent with the general notion of the speed-accuracy tradeoff.
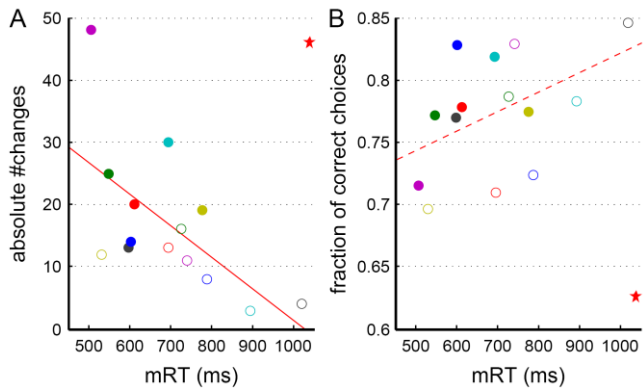
Fig. 5.6 Correlation between absolute number of changes, mean reaction time, and initial performance of individual participants.
(A) The overall number of changes is plotted against the mean reaction time (mRT) of all valid trials (2- and 4-choice condition together) for each participant (colored dots and circles). The red star denotes the one outlier subject that was excluded in the linear fits (red lines) and correlation analyses. (B) Relation of overall fraction of correct choices to mean reaction time. Individual subjects are marked as in (A). On average, participants with longer reaction times changed significantly less (A) and tended to be more accurate (B).

We also analyzed the correlation between the participants' overall accuracy and their ONC, but found no effect there (Pearson correlation coefficient $R = -0.23$, $p = 0.418$, N=14).

Taken together, as expected, faster reaction times led to lower performance, but also increased the probability to change the initial choice. We will return to the relation between changes of mind, reaction speed and accuracy in Section 5.4.2. Beforehand, we will describe the computational model and its fit to the subjects' average choice behavior in the next section.

5.4 Theoretical results

The experimental paradigm we used here to investigate changes of mind in the light of multiple alternatives is the combination of two recent experimental studies (Churchland et al., 2008; Resulaj et al., 2009). In the last two chapters, we showed that the results found in both of these studies can be explained by separate versions of a biophysically realistic decision-making model with attractor dynamics, based on the work of Brunel and Wang (2001; Wang, 2002). A critical question is whether the same theoretical concepts can now also account for the 2- and 4-choice RDM results we presented above.

Generally, attractor models implement the decision process by diffusion in a nonlinear landscape of stable fixed points, which act as decision-attractors. The decision alternatives correspond to sub-populations, or "pools", of excitatory neurons, which are selective for the respective motion directions (Fig. 5.2A, red). This means that they receive additional inputs if there is motion in their preferred direction. The output firing rates of these selective neural pools act as decision variables. Initially, they all fire at equal rates. With the onset of the motion input, which corresponds to the experimental RDM stimulus, the system dynamically evolves towards the decision state. In this network state, one of the selective pools fires at a high rate (winner), and the others are suppressed to low rates (losers). This "winner-take-all" competition arises through global inhibition, which is implemented by a population of inhibitory neurons connected to all neurons in the network.

The specific model we present here (Fig. 5.2), was developed by modifying the multiple-choice model with four selective pools presented in Chapter 4. As we have seen, with this model we already accounted for primate decision-making behavior and neural activity in an experimental paradigm with two and four alternatives, which was very similar to our 2- and 4-choice task, but without changes of mind (Churchland et al., 2008). Notably, all parameters in that model are independent of the number of choice alternatives. Solely the target input distinguishes between conditions in the model, in the same way as the visual R-targets determined the possible motion directions in the experiment. For two choice-alternatives the target input was applied to two selective populations, while for four alternatives, all four selective pools received the target input. Corresponding to the random-dot stimulus, the motion input was always applied to all four selective populations, independent of the number of alternatives (Eq. 5.2). A (first) decision was noted in the model, when one of the firing-rate transients crossed the decision threshold (38 Hz). If the firing rate of an initially losing selective pool subsequently exceeded the same decision threshold and surpassed the other pools by at least 5 Hz, that trial was considered a change of mind (see Fig. 5.2C for an example trial).

5.4.1 Model fit to average choice behavior

As can be appreciated in Fig. 5.7, the model matches the experimental reaction times and accuracy well for the different coherence levels. In comparison to the primate study (Churchland et al., 2008), the differences in reaction time between experimental conditions were less pronounced for our human participants and their accuracy was somewhat worse (see Fig. 4.2 for comparison). To adapt the model accordingly, we primarily changed three parameters: first, as we did not distinguish
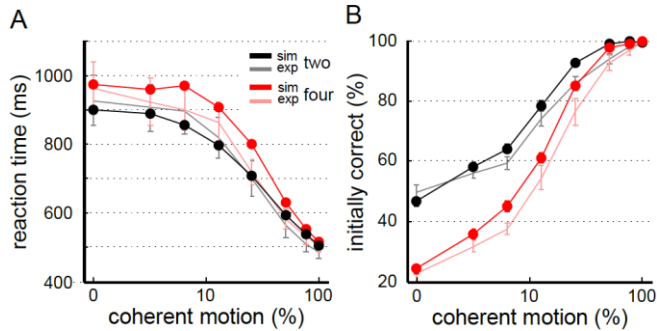
Fig. 5.7 Comparison between simulated and experimental reaction times and accuracy.
(A) Simulated RTs include a non-decision time (t_{ND}) of 350 ms. (B) Accuracy of the attractor model. For comparison the experimental data from Fig. 5.3 is shown again in lighter colors. Error bars of simulated data (SEM) are mostly hidden behind the dot markers.

whether two choices had a spatial distance of 180º or 90º, the connection strength between the selective pools now has no spatial component and is equal to ω_- for all selective pools. Second, the target input before motion onset (500 to 1,500 ms) was set much smaller, which lessened the difference between the 2- and 4-choice conditions. Third, the number of neurons in the network was reduced from 2,000 to 500. This mainly decreased the performance to a level comparable with our subjects' accuracy (Fig. 5.7B). Still, the model slightly overestimated the difference in reaction time between the 2- and 4-choice trials and the simulated performance slightly exceeded the experimentally observed accuracy. Nevertheless, these model adjustments can give insights into what might have caused the discrepancies between human and primate behavior. Roughly speaking, the model adaptations described above are congruent with less practice (see 5.5.3).

Most importantly for our purposes, the multiple choice model is able to produce changes of mind in the same way as the 2-alternative model (Chapter 3), which was used to fit the binary RDM experiment of (Resulaj et al., 2009), namely with a comparatively high motion input and a low decision threshold. Both, high inputs and a low decision threshold, generally lead to faster reaction times and less accuracy in the attractor model and thus correspond to pressure for speed in the experiment. Indeed, with a mean-field reduction of the full spiking-model, we have shown in Chapter 3 that this input-dependent speed-accuracy tradeoff arises from a shift of the dynamical working point in the attractor landscape of the network (Fig. 3.6). With higher inputs the system shifts closer to a bifurcation where a new attractor appears, which allows for

two neural populations firing at elevated rates. This means that after this bifurcation, decision-making is no longer unambiguous. The system there shows multi-stability between the decision attractors (corresponding to specific decisions) and the ambiguous, symmetric attractor (corresponding to two neural populations firing at elevated rates). Interestingly, as a consequence more changes of mind emerge the further the system is pushed towards this bifurcation. This is because in the proximity of the bifurcation, for high inputs, it becomes more likely that two selective populations reach firing rates close to the decision threshold, which facilitates changes of mind.

Importantly, this principle applies in the same way for the 4-alternative version of the attractor model. Notably, the 2-and 4-choice model predicts overall fewer changes for four choice-alternatives than for two. Thereby, it confirms our experimental results (Fig. 5.8, first column). Fewer changes of mind for more alternatives can theoretically be explained by the global inhibition in the network. If the target input is applied to four and not just two pools, inhibition and thus competition in
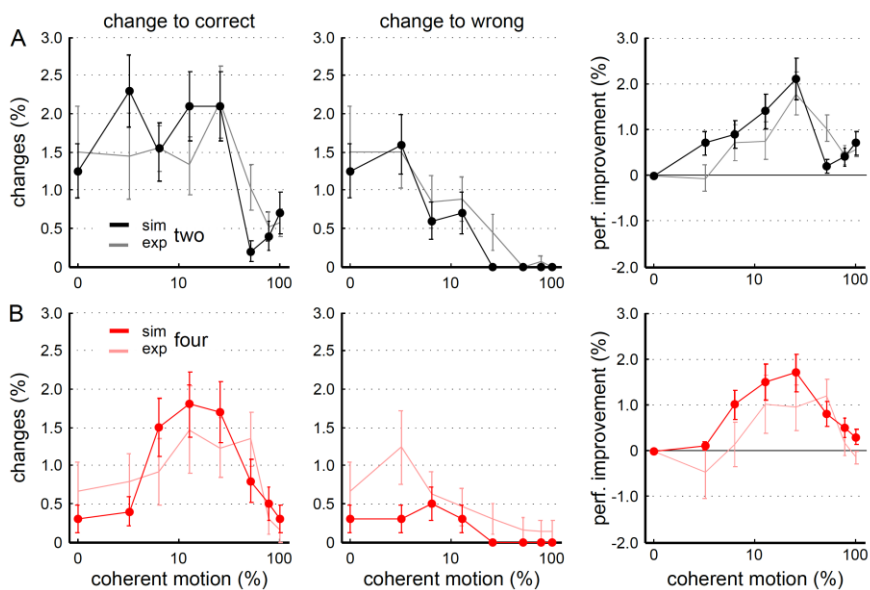
Fig. 5.8 Simulated changes of mind.
A model trial was counted as a change of mind, if, after the initial threshold crossing (first choice), the firing rate of an initially losing selective pool surpassed the other pools and crossed the decision threshold. (A,B, first two columns) Changes of mind as percentage of all valid trials for two (A) and four (B) alternatives. The experimental results of Fig. 5.4 and Fig. 5.5 are plotted in lighter colors for comparison.

the model increases. More competition has the opposite effect as higher inputs: it becomes less likely that two selective neural pools both reach firing rates close to the decision threshold, which is required for a change of mind.

What is more interesting is the fact that even the coherence dependence of the simulated changes matches that of the experimental changes of mind (Fig. 5.8). Also the relative difference between the 2- and 4-choice conditions in the performance improvement is captured by the computational model (Fig. 5.8, right column). Just for low coherences the absolute performance improvement is somewhat overestimated by the computational model. This is because for low coherences slightly more correcting changes were predicted by the model in the 2-choice trials, while in the 4-choice trials fewer erroneous changes occurred in the simulations than in the experiment.

In sum, our experimental results, including the changes of mind, are nevertheless in good accordance with the attractor model of decision-making and thus explained by the same theoretical concepts as the two preceding studies.

5.4.2 Participants grouped by ONC

Another benefit of the theoretical model is that it can help to elucidate the source of variability across individual subjects. In the following, we divided the 15 participants into three groups of 5 according to their readiness to change, i.e. their overall number of changes (ONC). Especially for low coherence levels, reaction times varied considerably between groups (see Fig. 5.9A and C for the 2- and 4-choice condition respectively). The group with most changes of mind responded most rapidly and vice versa, corresponding to the negative correlation between ONC and mean RTs that we found across individual participants (Fig. 5.6). Again similar to the analysis of individual subjects, hardly any trend was found in the grouped percentage of initially correct choices (Fig. 5.9, middle row). The overall percentage of correcting and erroneous changes for the three subgroups is displayed in the bottom row of Fig. 5.9 A,C.

Notably, in the computational model an adaptation of the decision threshold was enough to capture all subgroup differences described above (Fig. 5.9 B,D). The same threshold (38 Hz) that was used to fit the average across all participants also fitted the subgroup with an intermediate number of changes (Fig. 5.9, solid thin line). The simulated data for this group is thus identical to that shown in Fig. 5.7 and Fig. 5.8. Decreasing the threshold to 32 Hz produced faster RTs and more changes of mind in the 2- and 4-choice condition. Increasing the threshold to 44 Hz had the opposite effect. Moreover, this simple adjustment simultaneously fitted the quantitative differences of reaction times and

changes of mind between groups remarkably well. In addition, varying the decision threshold had only minor effects on the model's accuracy, in line with the experimental performance, which hardly differed between participant subgroups. Nevertheless, the group with most changes was the initially least accurate.

A further interesting detail can be noted in the coherence dependence of the changes of mind for the different subgroups (Fig. 5.A.2). Correcting changes in the model were particularly reduced for low coherences, if the decision threshold was set high. Interestingly, we observed the same effect in the experiment: The subgroup with least overall changes of mind did make correcting changes in the 4-choice condition, but only for high coherence levels, not for low motion coherence (Fig. 5.A.2, top right). In the model, this effect can be explained due to the nonlinear decision attractors. If the threshold is higher, it lies closer to the decision attractor. In the vicinity of the attractor, a first decision will only be reversed if the contrary evidence is strong enough to pull the transient out of the basin of attraction again. Therefore, initial errors will be corrected for high
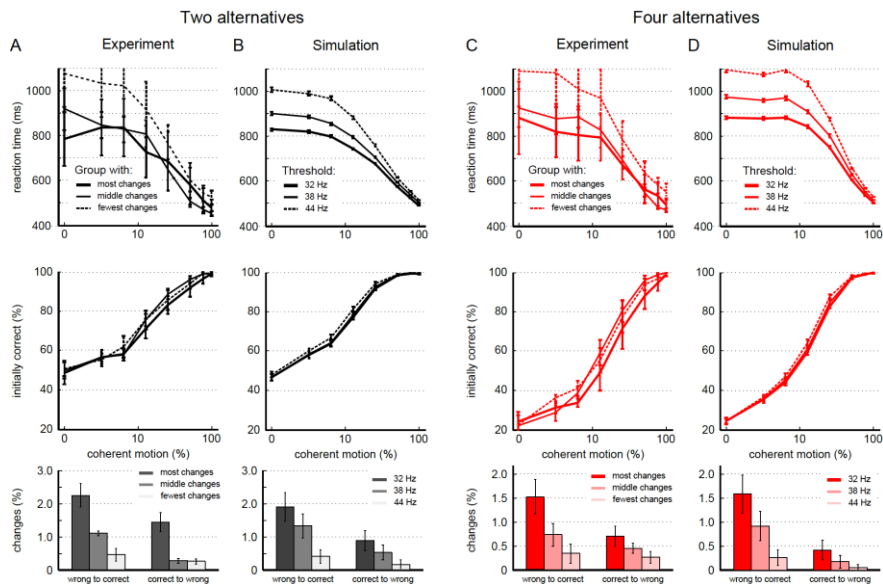
Fig. 5.9 Threshold variation accounts for differences in choice behavior of participants grouped according to their tendency to change.
(A,C) Experimental data. (B,D) Simulated data.(Top row) Reaction times. (Middle row) Initial performance. (Bottom row) Overall percentage of correcting and erroneous changes. Apart from the decision threshold, all other model parameters were kept constant. For a comparison of changes of mind for the different coherence levels see Fig. 5.A.2.

coherence levels even if the decision threshold is high, but not necessarily for lower coherences with weak evidence for a change of mind. This observation further affirms the connection between the decision threshold and the changes of mind we propose here.

Taken together, solely by an alteration of the decision threshold the nonlinear attractor model can explain the "changes-speed-accuracy" relation we observed experimentally in the behavioral variability of individual participants and subgroups.

5.5 Discussion

In this chapter, we presented human behavioral data from a 2- and 4-alternative random-dot motion discrimination task. Our participants reported their choice by moving a mouse pointer on the computer screen from a central start-target to one of the up-to-four response targets in the screen corners. In this setting, we were able to observe occasional changes of mind in the participants' movement trajectories. These changes of mind are supposedly based on information that was still unprocessed at the time of the first decision (Resulaj et al., 2009). A multi-alternative attractor model for decision-making that incorporates such further processing of evidence after the initial decision reproduced the experimental changes of mind and the general choice behavior of our participants.

In the following we will discuss our results with respect to the preceding studies that built the groundwork for our experiment. In particular, we will focus on our findings on differences between two and four alternatives, individual participants, and human versus primate choice behavior:

- Participants made more initial errors and less correcting changes for four than for two choices.

- The number of changes of individual participants negatively correlated with their reaction times.

- Despite behavioral differences, human and primate decision-making are explained by the same theoretical model.

- Furthermore, we will consider other concepts and modeling approaches, which might account for this set of results.

5.5.1 Comparison to binary changes of mind

Our findings on changes of mind in the 2- and 4-choice experiment are consistent with all relevant aspects of the original study on binary changes of mind (Resulaj et al., 2009). Changes leading to correct

responses were most frequent for intermediate coherences, erroneous changes decreased with higher coherent motion, and accuracy improved with changing. The present study thus showed that the general principles underlying changes of mind extend to multiple-choice decision-making.

Under closer inspection, however, in comparison to the three participants tested by Resulaj et al. (2009), our participants changed less often, had longer reaction times and showed much more between-subject variability. These quantitative behavioral differences can be explained by two alterations in the experimental setup: first, the 2- and 4-choice task was per se more difficult because of the different numbers of choice alternatives. Similarly, the dot-motion here ran along the diagonals, in direction of the R-targets and not simply to the left or right. Second, we simplified the experimental setup compared to (Resulaj et al., 2009). Instead of moving the handle of an elaborate vBOT manipulandum (Howard et al., 2009), our subjects used a standard computer mouse to indicate their choice.

To distinguish the effects of the simplified setup from actual task differences, we tested our subjects on the 2-top control (Fig. 5.A.1). There, they were presented with the same target configuration as in (Resulaj et al., 2009) and horizontal dot-motion. This led to ~250 ms faster RTs for low coherences and more correcting changes of mind compared to the 2-choice condition of the main experiment, which was randomly mixed with 4-choice trials (Fig. 5.A.1).

Interestingly, the participants' accuracy and erroneous changes of mind in the 2-top condition were basically identical to the 2-choice condition of the main experiment.

In the 2-top control our participants still changed on average less than those of Resulaj et al. (2009) and also the variability between subjects was higher. Both of these effects might result from less pressure to respond fast. Shifting the computer mouse involves only a very slight hand movement compared to moving a handle in space. The same time-out to start the motion response might thus have been less urging in our experimental protocol. Nevertheless, all of our participants responded substantially faster and less accurately than subjects which performed a similar 3-alternative RDM task and were free to respond without time limit (Niwa and Ditterich, 2008).

In sum, a variable number of alternatives with coherent dot-motion along the diagonals led to longer reaction times and less correcting changes of mind. Still, the basic principles of changes of mind extend to the case of multiple alternatives.

5.5.2 The "change-speed-accuracy" tradeoff

Due to the large number of participants we tested, we were able to evaluate the correlation between accuracy, reaction time, and changes of mind quantitatively (Fig. 5.6 and Fig. 5.9). Resulaj et al. (2009) already noted that their participants changed less when they were asked to respond more slowly. Indeed, we found a negative correlation between the overall number of changes and the mean reaction time across participants. This correlation was significant and thus even stronger than the negative trend we found between mRTs and the mean accuracy, which corresponds to the established concept of a speed-accuracy tradeoff. Grouping the participants by their overall number of changes revealed the same trends.

Notably, the theoretical model could explain the subgroup differences with different decision thresholds. The flexibility to reevaluate an initial choice might thus be regulated via the decision threshold. This is in line with our theoretical findings of Chapter 3, where we showed that changes of mind arise in the binary attractor model for high selective inputs and low decision thresholds. In theoretical models of decision-making threshold adaptation is usually associated with the speed-accuracy tradeoff (Palmer et al., 2005; Lo and Wang, 2006; Wong and Wang, 2006; Bogacz et al., 2010). Here, we added changes of mind to this relation: fast reaction times lead to more changes of mind, while changes of mind and accuracy seem to be barely linked with each other.

5.5.3 Comparison of human and primate choice behavior

In their groundbreaking study on multiple-choice decision-making, Churchland et al. (2008) measured the neural activity of macaque monkeys performing a multi-alternative RDM task. Here, we adopted their combined 2-and 4-choice paradigm, except that our participants were asked to indicate their choice by a continuous movement instead of a saccade, in order to observe changes of mind.

Interestingly, the primates showed substantial differences in reaction time for the different task conditions (Churchland et al., 2008). More precisely, while the monkeys' RTs were significantly slower in the 2-choice 90º than in the 2-choice 180º case, and slowest for four choices, we did not observe this dependence on the spatial angle for the human participants (compare Fig. 5.3 with Fig. 4.2). In fact, accuracy and RTs of the human subjects were similar for both 2-choice cases (90º and 180º), which is why we combined them here for comparison with the 4-choice condition. Moreover, also RT-differences between different numbers of alternatives were less pronounced for humans (The conditional differences in choice behavior are treated in more detail in Chapter 6).

In that respect, the computational model can help us to understand how the behavioral differences between the two species might arise, although decision-making in the RDM task should be based on the same principles. Notably, the model suggests, that training effects can largely explain the distinct behavior between species. This can be concluded from differences in crucial network parameters between the model implementation we used here, and the version that we applied to the monkey data in Chapter 4. There, we showed that a multi-alternative attractor model could account for all relevant aspects of Churchland et al.'s (2008) findings, including the differences in neural activity between two and four choice alternatives. The "monkey model" fitted the behavioral differences between the 2-choice 90º- and 180º-case observed for primates by assuming graded spatial connectivity: the connection weights between the selective neural populations that encode perpendicular motion directions were slightly stronger than between those selective pools that encode opposite motion directions. Monkeys are usually trained for months on a psychophysical task, before the final experiment is conducted. Such a spatial connectivity component might thus emerge during training through Hebbian learning, assuming that neural populations that encode more similar motion directions have been firing more correlated in the past.

In contrast, to fit the behavior of the untrained human participants we had to drop this spatial connectivity component. Moreover, to approximate the somewhat worse performance of our human subjects, we reduced the network size from 2,000 to 500 neurons, which increased finite size noise in the model and thereby reduced its accuracy. With training, more neurons might be recruited to encode possible motion directions.

Overall, the comparison of the computational models revealed that human and primate decision-making can be accounted for by the same theoretical mechanisms. Accordingly, our results generally agree with Churchland et al.'s (2008) findings for primates, despite deviating behavioral observations. What is more, the theoretical model provides a testable prediction on primate behavior regarding changes of mind. The monkey-model (Chapter 4) would predict more changes of mind between R-targets at a spatial angle of 90º than for 180º, due to the stronger connection between the selective pools, and even fewer changes for four than for two choices. It would be interesting to see if the overall frequency of changes of mind in trained monkeys, or overtrained humans, indeed varies more between the different experimental conditions.

5.5.4 Intuition and possible models for changes of mind

One might have suspected that the overall probability to make a change of mind would be higher for four alternatives because of more confusion between more possible motion directions. Added to this, the prior probability to choose a particular direction without further evidence is 50% for two and 25% for four alternatives, if all directions are assumed equal. Yet, while accuracy was indeed lower for four alternatives according to the different prior probabilities, changes of mind seem to occur more often the easier the task is.

Moreover, the different a priori probabilities make it difficult to form an intuition about the dependency of correcting and erroneous changes on the number of alternatives. Strikingly, the probability to discard an initially correct R-target seems to depend only on the given level of evidence, which here means the coherence level in the motion stimulus: Erroneous changes of mind occurred independent of the number of alternatives and target locations relative to the motion direction (Fig. 5.4D and Fig 5.A.1G). By contrast, the probability to correct an initial error was much higher for two alternatives than for four, and even higher for easier left/right choices (Fig. 5.4C and Fig. 5.A.1E, respectively). Even if the prior probabilities are taken into account, there is no simple intuitive explanation for these findings.

Notably, the attractor model of decision-making produced the changing behavior of the 2- and 4-choice experiment through its global inhibition, while the threshold and all other network parameters were identical across conditions.

Attractor models are however not the only decision-making models that are generally capable of producing changes of mind. Resulaj et al. (2009) extended a linear diffusion model by an independent second threshold that determined changes of mind. This model could then be fitted to their experimental findings on changes between two alternatives (left/right). Because in the diffusion model it is the difference in evidence between two choices that is integrated, an extension to more than two alternatives is not straightforward. Yet, Churchland et al. (2008) could account for their 2- and 4-choice data by a race between two independent diffusion models and (Niwa and Ditterich, 2008) implemented a diffusion process between three alternatives through weighted feedforward inhibition. It is probable, that a similar 4-alternative version of these diffusion models could also reproduce our experimental observations, especially, if the threshold for changes of mind were a free parameter as in (Resulaj et al., 2009). Note again, that in the physiologically-inspired attractor model the change threshold is the same as the decision threshold. It would be very interesting to see how our experimental results on changes of mind between multiple alternatives would constrain the above

and other existing multiple-choice decision models (Beck et al., 2008; Furman and Wang, 2008).

Taken together, we have discovered that more alternatives lead to more initial errors and a lower probability to improve the initial choice. Through the combined approach of experiment and theoretical model, we were able to further establish and explain a relation between changes of mind, reaction speed, and accuracy over the decision threshold. Changes of mind in multiple-choice decision-making might further help to distinguish between theoretical models in the future.

5.A Chapter appendix

5.A.1 Choice behavior in the 2-top control condition

In a separate block of trials we replicated the experimental design of Resulaj et al. (2009) with two R-targets at the top of the screen and the starting point at the bottom ("2-top" condition) with our simpler setup using a computer mouse instead of a handle.

While the reaction times in the 2-top condition were much faster, than for two alternatives in the 2- and 4-choice paradigm, performance is basically identical. Note that in the 2-top condition, the possible motion directions were always horizontal (left or right), while in the main experiment coherent dots moved along the diagonals. Correcting changes were more frequent in the 2-top condition than for two alternatives in the main experiment. Erroneous changes, however, happened with the same frequency. Thus, the performance improvement was greater in the 2-top condition.

5.A.2 Model with adapted thresholds matched frequency distributions of changes for participants grouped by ONC

A threshold of 32 Hz was used to fit the choice behavior of the group with most changes and 44 Hz for the group with least changes. For the middle group the same threshold (38 Hz) was used as in Fig. 7 and 8 in the main text to match the average across all subjects. All other network parameters were kept constant. Note that in the attractor model the number of correcting changes (B,D, top row) decreased by increasing the threshold, especially for low coherence values. The change distributions thus became flatter. That is because, as the threshold approaches the nonlinear decision attractor, a first decision will only be reversed if the contrary evidence is strong enough to pull the transient out of the basin of

attraction again. Interestingly, we observed the same effect in the experiment (A,C, top row).
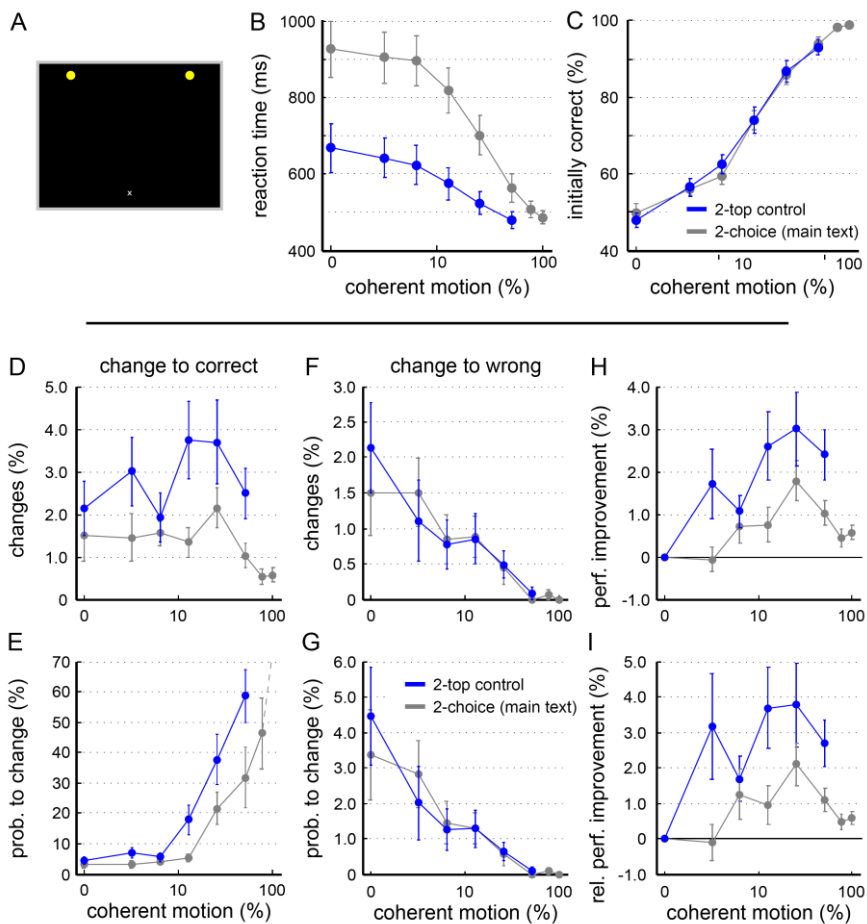
Fig. 5.A.1 2-top control condition.

(A) Target configuration. (B) Reaction times and (C) performance averaged across all 15 participants. The 2-alternative condition from the main text is plotted in gray for comparison. (D-G) As Fig. 5.4. (D,F) Changes of mind as percentage of all valid trials. (E,G) Probability to change, i.e. the same data as in (D,E) but divided by the number of initially correct or error trials the changes could have originated from. (H,I) as in Fig. 5.5. (H) Absolute difference of initial and final performance (considering changes of mind). (I) Relative performance difference, i.e. the performance difference shown in (H) divided by the initial performance.
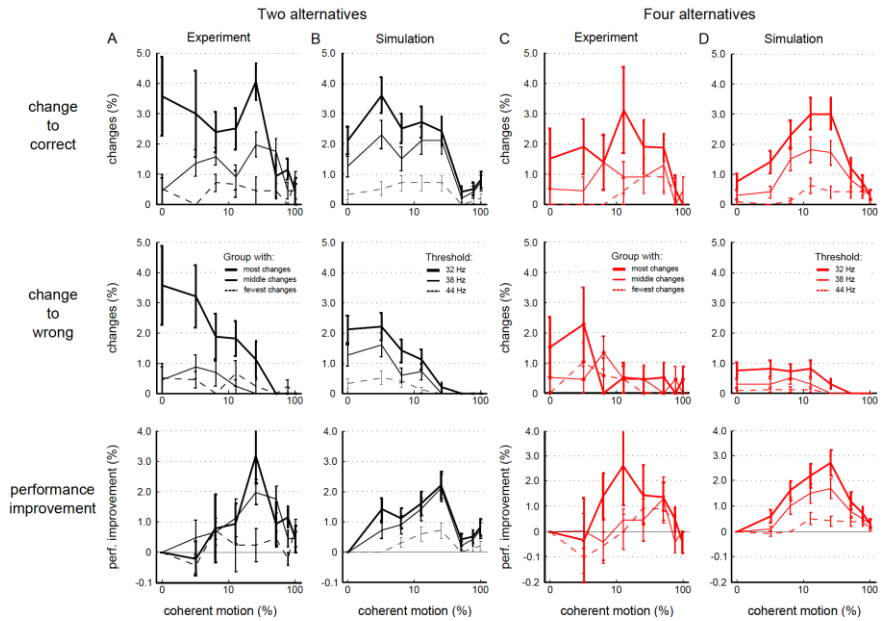
Fig. 5.A.2 Threshold adaptation explains distribution of changes for different participant groups.

(A,C) Frequency of correcting (top row) and erroneous (middle row) changes of participants grouped according to their overall number of changes (ONC) as percentage of all valid trials. (B,D) Simulated change distributions for three different decision thresholds.

6 A VISUAL ILLUSION IN THE RDM-TASK?

6.1 Introduction

In the last chapter, we have presented our experimental and theoretical results on the question how changes of mind depend on different numbers of choice-alternatives. We approached this topic in a 2- and 4-choice random-dot motion discrimination task, in which human subjects had to indicate their choice by moving a computer mouse to the selected response target on the screen. In fact, our experiment combined the two RDM protocols used by Resulaj et al. (2009) and Churchland et al. (2008) to study changes of mind and multiple-choice decision-making independently (Chapter 3 and 4).

In the RDM task, more choice-alternatives automatically imply a smaller spatial angle between the possible directions of coherent motion. According to Churchland et al. (2008), for an accurate comparison of two and four alternatives, we thus included a third experimental condition with two alternatives at an angular distance of 90º, in addition to the standard 2-choice condition with 180º angular distance, and the 4-choice condition. The decision behavior of our human participants resulted quite similar in the 2-choice 90º- and 180º-cases. As we have discussed in Section 5.5.3, this is contrary to the decision behavior observed in monkeys performing the 2- and 4-choice RDM task, which needed significantly longer to decide between 2-choice targets at the smaller angle of 90º (Churchland et al., 2008). To filter the actual effects of the different numbers of alternatives from influences of the different spatial angles, we thus collapsed the 2-choice 90º-case and the 2-choice 180º-case, with the additional benefit of a greater sample size for comparison to the 4-choice case.

Here, we elaborate on our experimental results, considering the different experimental conditions in further detail. We first deal with a small but significant difference in performance between the two conditions with two choice-alternatives. With a slight modification of the network inputs to the attractor model of Chapter 5, we could account for this directional difference of 2-choice performance. In particular, we introduced heterogeneity in the motion bias that would correspond to a visual illusion of reversed coherent motion in the perception of the RDM stimulus. Based on this "illusion bias", we could further explain some initially puzzling observations in the 4-choice condition, regarding the direction of errors and changes of mind with respect to the direction of coherent motion. More precisely, we found that errors and changes of mind to the R-target 180º opposed to the correct R-target were relatively

more frequent than expected by chance. On the one hand the findings presented here affirm the explanatory power of the biophysically realistic attractor model. On the other hand, they send a note of caution with regard to the precision of the perceived motion coherence in RDM experiments. Nevertheless, our main experimental conclusions on how changes of mind depend on the number of choice-alternatives (Chapter 5) are not subject to the apparent visual illusion.

The main results of this chapter are:

- The attractor model can account for performance differences in the 2-choice 90º- and 180º-cases assuming an illusion bias in the motion input.

- The illusion bias also elucidates the disproportional frequency of 180º-errors and changes of mind in the 4-choice condition.

- The assumed illusion does not affect the overall dependence of changes of mind on the number of alternatives.

6.2 Methods

The experimental data set we describe and discuss in this chapter is the same as that presented in Chapter 5. In a 2- and 4-choice RDM paradigm, we tested 15 human subjects in three different experimental conditions, regarding the number of choice-alternatives and the relative position of the response targets, which indicate the possible directions of coherent motion. The participants were either presented with two choice-alternatives at an angular distance of 180º, two alternatives at an angle of 90º, or four choice-alternatives that were equally distributed with 90º angular distance (Fig. 5.1B). The experimental paradigm is described in detail in Section 5.2.1 and Appendix A.3.1.

Also the computational model we employ here, to account for our experimental findings is identical to the biophysically realistic attractor model of Chapter 5, endowed to simulate changes of mind in a decision-making task with up to four choice-alternatives. Its general characteristics are summarized in Fig. 5.2 and Section 5.2.2. More details are given in the Appendix A.1 and Table A.5. All default simulation parameters are listed in Table A.6.

The only alteration between the model simulations we report here and that of Chapter 5 concerns the motion input applied to the selective populations, which corresponds to the experimental RDM stimulus. In order to simulate the effects of the proclaimed visual illusion of reversed coherent motion, we introduced a heterogeneity in the negative input bias applied to the unfavored selective pools. The total motion input to the

network is still the same for all coherence levels, but instead of Eq. 5.2, we used the following motion input:

$$\nu_{\text{motion},1} = 40 \text{ Hz} + 30 \text{ Hz} \left(1 + \frac{3 \cdot coh}{100\%} \right),$$

$$\nu_{\text{motion},2/4} = 40 \text{ Hz} + 30 \text{ Hz} \left(1 - \frac{(3 - \gamma)/2 \cdot coh}{100\%} \right), \text{ and} \qquad (6.1)$$

$$\nu_{\text{motion},3} = 40 \text{ Hz} + 30 \text{ Hz} \left(1 - \frac{\gamma \cdot coh}{100\%} \right).$$

The "illusion factor" γ that was added to the negative bias of selective pools 2-4 determines the inequality in motion bias received by selective pool 3 compared to pools 2 and 4. In the following we used either $\gamma = 0.2$ or $\gamma = 0.4$. Note that $\gamma = 1$ would correspond to the original motion input used in Chapter 5 (Eq. 5.2). For $\gamma < 1$, selective pool 3, with direction preference opposite to the direction of coherent motion, receives a less negative bias than selective pools 2 and 4, whose direction preference is perpendicular to the coherent motion. In this way, we implemented our hypothesis that the net direction of motion perceived by the participants could sometimes be opposite of the actual direction of coherent motion, due to a visual illusion. As a consequence of the illusion factor γ, the maximal difference between the sensory inputs (target and motion input) to the selective pools now depends on the R-target locations (Fig. 6.2). The implications of this heterogeneous motion bias on the simulated decision behavior will be presented in the following.

For statistical tests, we used within subject ANOVAs on the initial performance and reaction times of all correct trials, excluding changes of mind, with the factors "angular distance" (2-choice 90º vs. 2-choice 180º) and "coherence" (eight levels of coherence: 0%, 3.2%, 6.4%, 12.8%, 25.6%, 51.2%, 76.8% and 100%).

6.3 Results

6.3.1 Behavioral differences in the 2-choice 90º- and 180º-cases

The mean reaction times and initial performances of the two 2-choice conditions from our 2- and 4-choice RDM discrimination task are displayed separately in the left column of Fig. 6.1. Participants responded somewhat more accurately in the 90º-case, while their mean reaction times did not differ substantially for the two 2-choice conditions. Statistical tests on reaction times and the percentage of correct responses revealed the main effects of "angular distance" (F(1,14) =9.862, p < .007)

and "performance" (F(1,14) =8.309, p < .012). A second level analysis showed that the performance in the 90°-case was specifically better than in the 180°-case for the following intermediate coherence levels: 12.8% (p < .001), 25.6% (p < .021) and 51.2% (p < .039). The mean RTs were not significantly different between the 90°- and the 180°-cases, except for 0% motion coherence (p < .003).

Curiously, the better performance and equal mean RTs would suggest that the 2-choice 90°-case was actually easier for our human subjects, contrary to Churchland et al.'s (2008) findings for monkeys, which needed longer to respond in the 90°-case with the same level of accuracy as in the 180°-case (compare Fig. 4.2B,D).

In line with this inconsistency, we were not able to reproduce the better performance in the 90°-case with the 2- and 4-choice attractor model, while at the same time accounting for similar reaction times and changes of mind in both 2-choice conditions. A first hint for an alternative explanation of the better 90°-case performance came from the comments
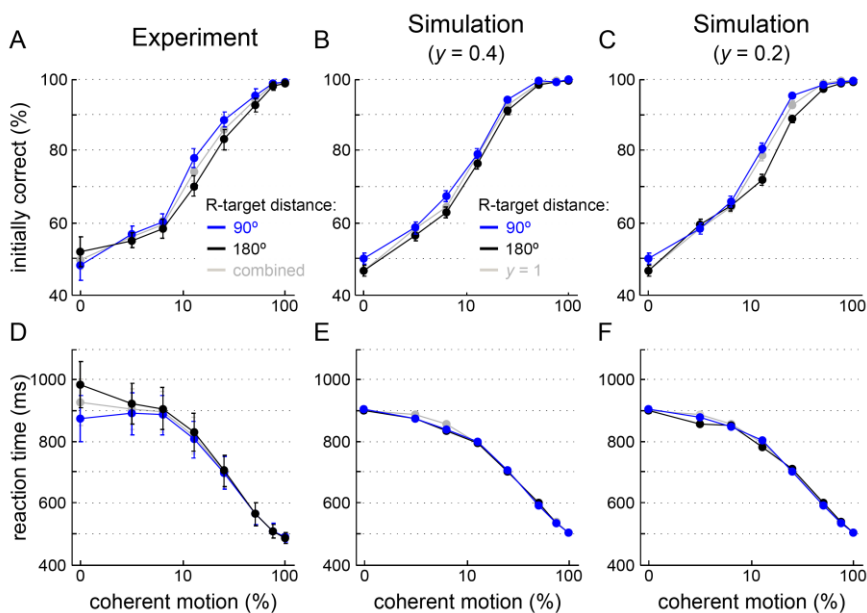
Fig. 6.1 Performance and mean RTs for the 2-choice 90°- and 180°-case.
(A,D) Experimental Data. For intermediate motion coherence the initial performance was better in the 90°-case. Mean RTs did not differ significantly for the two conditions. (B-D) Model simulations with illusion factor γ = 0.2 or γ = 0.4. With increasing heterogeneity in the motion bias (smaller γ), the performance of the 180°-case becomes worse, while mean RTs are basically unaffected. The combined 2-choice case and the model simulations of Chapter 5 (γ = 1) are plotted in gray for comparison.

of a few participants after performing the experiment: When asked if they observed anything remarkable during the experiment, some participants stated that occasionally they clearly perceived coherent motion in the direction opposite to the correct direction.

Based on the participants' observations, we hypothesize that a visual illusion of reversed coherent motion might cause the performance difference in the 2-choice conditions. In particular, we assume that, at the onset of coherent motion, especially for intermediate coherence levels, one might get the illusory impression that the coherent dots were stationary, while the randomly moving dots seem to shift into the opposite direction. Moreover, this illusion might affect responses more in the 180°-case than in the 90°-case, as the direction opposed to coherent motion is not a valid option in the 90°-case (there is no response target located in the illusionary direction).

We tested our hypothesis with the biophysically realistic attractor model of Chapter 5, which accounts for choice behavior including changes of mind for two and four choice-alternatives, if the 90°- and 180°-cases are combined. To implement the visual illusion in the model simulations, we introduced an "illusion factor" γ that produced an inequality in the negative motion bias (Eq. 6.1). Previously, we subtracted the same motion bias ($coh\cdot30$ Hz) from each of the three "unfavored" selective pools (selective pools 2-4; preferred direction of selective pool 1 was set to correspond to the direction of coherent motion). In the context of the illusion factor, a uniform negative bias equates to $\gamma = 1$.

In Fig. 6.1B-D the simulated initial performance and mean RTs of the attractor model with illusion bias are shown for $\gamma = 0.4$ and $\gamma = 0.2$. Indeed, with decreasing γ the performance in the 180°-case deteriorates compared to the 90°-case. Furthermore, as in the experimental data, mean RTs are still basically identical for the 90°- and 180°-case. Reaction time is thus not affected by the illusion bias.

What causes this conditional difference in performance in the computational model? In Fig. 6.2 we show an illustrative example of the target inputs and the motion bias for $\gamma = 0.4$. An illusion factor of $\gamma < 1$ effectively leads to a less negative motion bias applied to selective pool 3, which encodes motion that is 180° opposed to the direction of coherent motion. The perpendicular selective pools 2 and 4, on the other hand, receive even more negative input biases. Through the respective target inputs, the relative location of the R-targets now determines the difference between the total sensory input received by selective pool 1 and the second highest input to one of the other selective pools. In the 180°-case, the input difference between the correct alternative and the second alternative is smaller than in the 90°-case. To make this more clear, in Fig. 6.2, below the respective selective pools, we indicated the sum of the motion bias plus target input for 100% coherent motion. The smaller
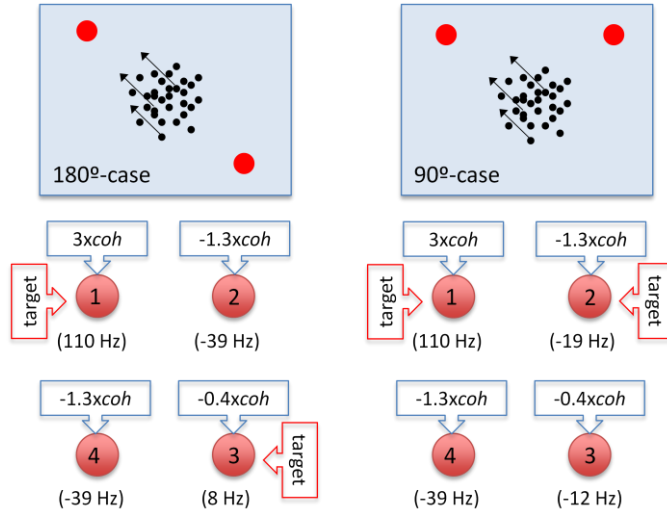
Fig. 6.2 Input differences in the 2-choice 90º- and 180º-case for $\gamma = 0.4$.
Red, numbered circles depict the four selective pools. The red and blue arrow-boxes indicate the target input and motion bias for the case of coherent motion into the preferred direction of selective pool 1. In brackets below the respective selective pools, the absolute amount of target input + motion bias is given for 100% motion coherence.

difference in total sensory input to the two choice-alternatives explains why performance in the 180º-case is worse than in the 90º-case.

In sum, the illusion bias together with the target input gives rise to conditional differences, based on different total sensory inputs to the choice alternatives. By simulating a visual illusion of reversed motion coherence in the RDM stimulus we could thus account for the experimentally observed performance difference between the 90º- and 180º-case.

The left column of Fig. 6.3 shows the changes of mind for two choice-alternatives, separated into the 90º- and 180º-case. Overall, the frequency of changes of mind with respect to motion coherence did not differ substantially for the two 2-choice conditions, neither for correcting changes (from an initially wrong choice to the correct R-target), nor for erroneous changes (from the correct to the wrong choice). Consequently, averaging across the 90º- and 180º-case, as we have done in Chapter 5, neither has distorted the overall frequency of changes of mind, nor their dependence on the amount of coherent motion. Only for 25.6% motion coherence we observed more correcting changes of mind in the 180º-case.

In line with this, the attractor model with illusion bias ($\gamma = 0.4$) produced somewhat more correcting changes of mind in the 180º-case

than in the 90°-case (Fig. 6.3, right column). This effect again results from the smaller difference in total input between selective pool 1 and pool 3 in the 180°-case compared to the difference between pool 1 and 2 in the 90°-case. If the input difference between selective pools is smaller, a change becomes more likely. Changes from the correct to the wrong alternative occurred with about equal frequency in both 2-choice conditions.

Moreover, it has to be noted that the simulated frequency of changes of mind from Chapter 5, without illusion bias ($\gamma = 1$), was mostly intermediate between the 90°- and 180°-case for $\gamma = 0.4$. Consequently, at least in the model, the simulated visual illusion did not affect the overall probability of changes of mind for two choice-alternatives.
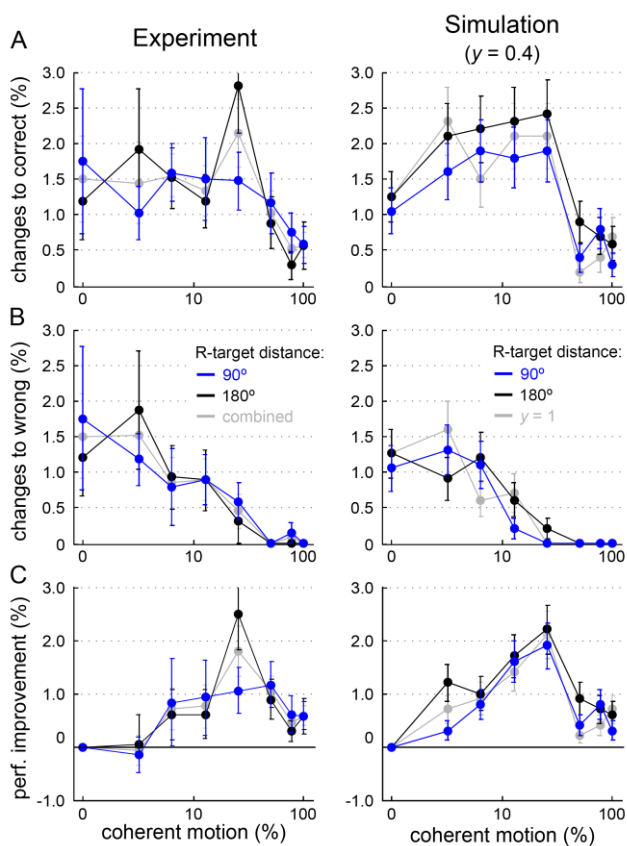
Fig. 6.3 Comparison between changes of mind in the 2-choice 90°- and 180°-case.

(A) Correcting changes (B) and erroneous changes are shown as percentage of all valid trials for experimental data (left column) and model simulation (right column). (C) Performance improvement with changes of mind (absolute difference of initial and final performance).

121

Before turning to the 4-choice condition, we briefly return to the peak in correcting changes of mind for the 180º-case at 25.6% coherent motion. The computational model gave one explanation for a higher number of correcting changes in the 180º-case. Another (additional) possibility would be that the visual illusion might vanish shortly after motion onset, leading to a change of the perceived motion direction from the illusionary to the actual direction of coherent motion, which could cause more changes of mind.

6.3.2 A closer look at the 4-choice condition

Perhaps the most convincing experimental evidence for the existence of the hypothesized visual illusion can be found in the direction of errors in the 4-choice condition. In the RDM stimulus, the dots that are not moving coherently are moving around randomly. Hence, in the 4-choice condition, errors should be equally distributed across the three incorrect R-targets. The chance probability for an error trial, in which a R-target perpendicular to the correct choice was selected, should thus be 67%. Accordingly, errors to the R-target opposed to the correct choice should constitute 33% of the total number of errors. The attractor model without illusion bias ($\gamma = 1$) naturally exhibits this chance distribution of errors, as all distracters (selective pools 2-4) receive the same amount of sensory inputs and are also otherwise identical (Fig. 6.4, dashed line).

The thin solid lines in Fig. 6.4 depict the actual, experimentally observed fractions of perpendicular and diametrical errors ("90º- and 180º-errors"). For small coherence levels ($\leq 6.4\%$) errors are still approximately partitioned according to chance. At 25.6% motion coherence, 180º-errors occurred with slightly higher frequency as 90º-errors. For even higher coherence levels probabilities for 90º- and 180º-errors were actually reversed.

The attractor model with illusion bias ($\gamma = 0.4$) could reproduce the experimental dependence of error direction on motion coherence remarkably well (Fig. 6.4, thick solid lines). Decreasing the illusion factor further, would lead to a reversal of the probabilities for 90º- and 180º-errors already at lower motion coherence (not shown).

Analogous to the directionality of errors, also changes of mind in the 4-choice condition can either be directed to a perpendicular, or to the opposite R-target, with respect to the initial choice (90º- and 180º-changes). Fig. 6.5 shows the experimental and simulated percentage of 90º- and 180º-changes of mind for four choice-alternatives. Note that a correcting 90º-change must have originated from an initial 90º-error. In that way, the proportion of correcting 90º- and 180º-changes of mind directly depends on the respective proportion of 90º- and 180º-error trials. Accordingly, for intermediate coherence levels correcting 180º-changes
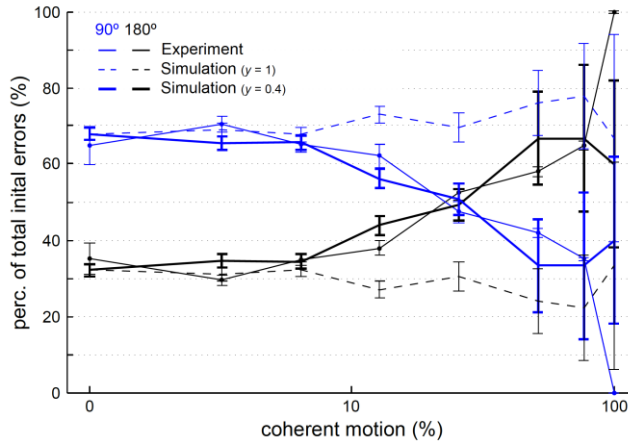
Fig. 6.4 Direction of errors relative to the correct choice in the 4-choice case.
Thin solid lines denote the experimental percentages of error trials directed to a R-target with 90° (blue) or 180° (black) angular distance from the correct R-target. They are at chance level (67% and 33%) for low coherences, but become reversed for high coherence levels. The model with illusion motion bias replicates these experimental findings (bold solid lines), while for uniform negative biases (γ = 1, dashed lines) the fractions of 90°- and 180°-errors naturally stay at chance level. Note that error bars increase for higher coherence as the total number of errors decreases.

occurred more frequently than would be predicted by chance (Fig. 6.5A, left panel). As for the errors, this behavior is captured by the attractor model with illusion bias ($\gamma = 0.4$), but not without the illusion bias ($\gamma = 1$) (Fig. 6.5A, middle and right panel). Moreover, the same logic as for error trials applies to erroneous changes of mind. By chance, 67% of changes from the correct to a wrong target should be 90°-changes, 33% 180°-changes (see simulation with $\gamma = 1$; Fig. 6.5B, right panel). While for high coherences all erroneous changes are indeed 180°-changes, congruent with the attractor model with illusion bias, we strangely did not observe the same effect at intermediate coherence levels.

For wrong-to-wrong changes illusion effects are hard to detect, as both with and without the hypothetical illusion most wrong-to-wrong changes should be 90°-changes. Without the visual illusion this is because of chance, with the illusion there are more initial 180°-errors which relatively increase the number of 90° wrong-to-wrong changes.

Overall, as for the 2-choice case, the total percentage of change-of-mind trials and its dependence on motion coherence in the model were not affected by the simulated visual illusion. Only the relative proportion of 90°- and 180°-changes were altered by the illusion bias.
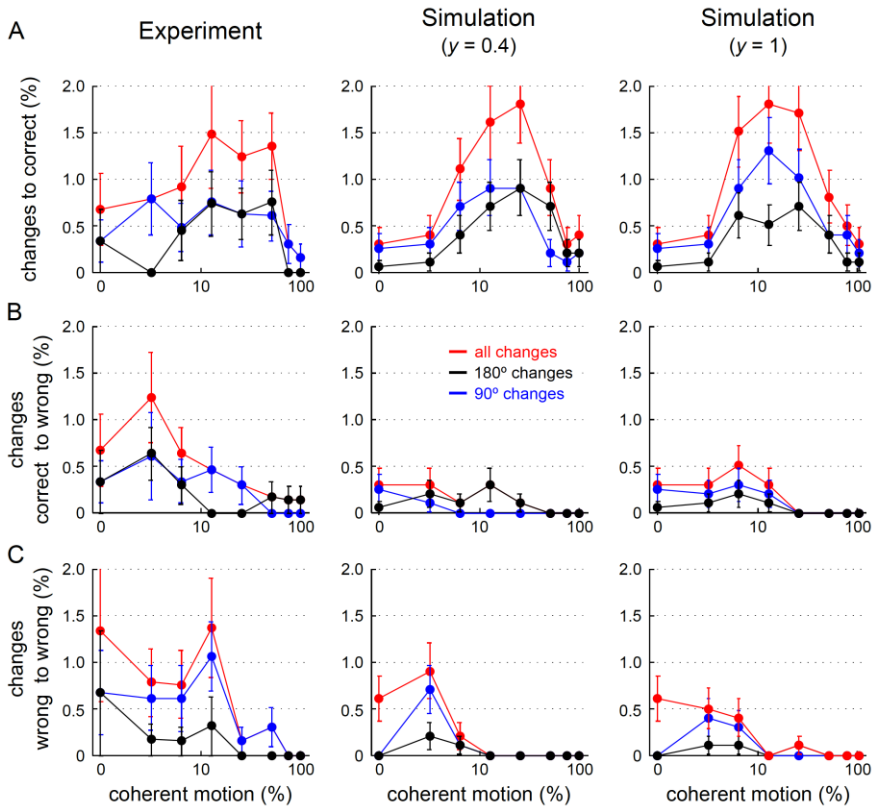
Fig. 6.5 90º- and 180º-changes of mind in the 4-choice condition.
(A) Correcting, (B) erroneous, and (C) wrong to wrong changes are shown as percentage of all valid trials for experimental data (left column) and model simulations with γ = 0.4 and γ = 1 (middle and right column). The total percentage of changes is displayed in red. For experiment and model simulation with γ = 1 it is the same as shown in Fig. 5.4 and Fig. 5.8, respectively.

Despite the amount of initially puzzling observations that could be resolved by incorporating the visual illusion, a few points in the experimental data remain unclear. In particular, the peak in the percentage of wrong-to-wrong changes of mind for 12.8% coherent motion could not be explained by the model. In fact, the frequency of wrong-to-wrong changes of mind would be expected to decrease monotonically with increasing coherence, in the same way as for erroneous changes. This is because large random fluctuations in the RDM stimulus, which might cause such changes, also become less likely with increasing coherence.

Moreover, the model still predicts to few erroneous and wrong-to-wrong changes in the 4-choice condition. One way to increase the number

of changes in the 4-choice relative to the 2-choice condition might be to add a top-down input dependent on the number of alternatives (Furman and Wang, 2008).

Nevertheless, apart from these minor issues, the biophysically realistic attractor model captures all relevant aspects of our experimental data. Indeed, the fact that we could explain unexpected experimental observations by simulating an assumed visual illusion, convincingly affirms the general validity of the attractor model.

6.4 Discussion

In this chapter, we continued the account of our psychophysical 2- and 4-choice RDM experiment with changes of mind, which we introduced in Chapter 5. While the last chapter was dedicated to experimental observations directly relevant to perceptual decision-making, here we focused on more subtle aspects of our 2- and 4-choice RDM data set. In particular, we treated the two 2-choice conditions, the 90º- and 180º-case, separately, and also considered errors and changes of mind in the 4-choice condition with respect to their directionality (90º or 180º from correct choice or initial decision). We further showed that several minor observations that could not be accounted for by the 2- and 4-choice attractor model described in Chapter 5 could eventually be fit if an "illusion bias" was added to the sensory motion input. Thereby, the attractor model affirmed our hypothesis that a visual illusion of reversed coherent motion might be responsible for the observed discrepancies in our experimental data.

In the following, we will discuss the implications of the proclaimed visual illusion in the random-dot motion stimulus with respect to the results presented in Chapter 5 and preceding studies on visual-motion discrimination. Beforehand, in the next section we briefly address the implementation of the hypothesized illusion in the attractor model.

6.4.1 Simulating the visual illusion

A few of our participants reported that on a small fraction of trials they actually perceived coherent motion towards the response target 180º apart from the correct response target. This illusionary impression might arise at the onset of coherent motion for intermediate coherence levels, if the coherently moving dots were perceived as stationary. In that case, the randomly moving dots would seem to shift jointly into the opposite direction as the coherent dots.

With the computational model we aimed to assess the influence of this illusionary impression on our behavioral findings. In particular, we were interested whether the assumed visual illusion might resolve

remaining discrepancies between our model simulations of Chapter 5 and several minor peculiarities in the decision behavior of our participants.

To simulate the visual illusion, we introduced an "illusion bias" in the model input that effectively increased the sensory evidence for the direction 180º opposed to coherent motion, while it decreased the evidence in favor of the perpendicular directions[18]. The strength of the illusion bias in the model was regulated by the "illusion factor"γ, which was arbitrarily adjusted to match our experimental observations. For simplicity we used the same γ for all coherence levels. Furthermore, all particular characteristics of the proclaimed visual illusion, such as the percentage of trials that are affected by the illusion, are combined in this illusion factor and not further explored here.

Regardless of our simplifying assumptions, by adding the illusion bias to the model, we could account for most of the minor behavioral discrepancies. Indeed, despite our efforts and the large number of parameters, we did not succeed to fit the model without illusion bias to the difference in performance between the 2-choice 90º- and 180º-cases, in addition to the remaining decision behavior. This fact further shows that the attractor model cannot simply reproduce any virtual decision behavior, as it is restricted by its physiological construct. As a consequence the model can postulate additional, initially unregarded influences, such as the proclaimed visual illusion.

6.4.2 Implications of the visual illusion on the validity of our behavioral results

In Chapter 5, we have presented our findings on changes of mind during multiple-choice decision-making based on the 2- and 4-choice RDM discrimination task. The results of the current chapter, however, suggest that the RDM stimulus used in this task lead to a visual illusion which affected the perceived amount of coherent motion. Even if only a small portion of trials was subject to this visual illusion, it nevertheless raises the question, whether the conclusions drawn in Chapter 5 can be sustained in the light of this experimental contamination.

Here, the theoretical model provides reassuring evidence that the assumed visual illusion did not affect the main experimental conclusions presented in Chapter 5. A moderate change in the model's input bias was

---

[18] It has to be noted that in this way, we did not actually implement the visual illusion itself, but rather its perceptual effect of reversed coherent motion. In other words, we assume that sensory activity in area MT might be altered by the visual illusion. As before, the attractor network represents a local microcircuit in area LIP, which receives sensory inputs from area MT that is subject to the illusionary percept.

sufficient to capture most of the previously unexplained observations with the attractor model. This already indicates that the assumed visual illusion could not have dramatic effects on the overall decision behavior. In fact, simulated reaction times were basically unaffected by the illusion bias (Fig. 6.1). Furthermore, the overall number of changes for two and four choice-alternatives did not depend on the illusion factor. More precisely, although the illusion bias in the model lead to somewhat more changes of mind in the 2-choice 180º-case compared to the 90º-case, the average percentage of changes across the two conditions was the same for $\gamma = 0.4$ and $\gamma = 1$.

Moreover, the attractor model suggests that the experimentally observed differences in performance between the 2-choice 90º- and 180º-cases can be attributed to the visual illusion. This further affirms our conclusion of Section 5.5.3 that the decision behavior of our naive human participants was indeed not affected by the different angular distances between the R-targets in the 2-choice 90º- and 180º-case, contrary to the behavior of trained monkeys (Churchland et al., 2008).

In all, with respect to changes of mind, the illusion bias only influenced the relative proportions of 90º- and 180º-changes, but not the total amount of changes for the different numbers of alternatives. Therefore, we conclude that none of the results presented in Chapter 5 was distorted by the visual illusion.

6.4.3 Could previous experiments have been influenced by the illusion?

In the last section we have argued that our findings on changes of mind in relation to the number of choice alternatives are almost certainly not affected by the proclaimed visual illusion. Yet, it had an influence on other more detailed aspects of our data, like the performance in the two 2-choice conditions. As the RDM task is a widely used paradigm, one has to ask whether previous studies might also have been subject to this visual illusion.

In that regard, it first has to be said that the visual stimuli we used in our experiment were very similar to preceding studies (Palmer et al., 2005; Churchland et al., 2008; Niwa and Ditterich, 2008; Resulaj et al., 2009). We used standard parameters for the motion stimulus (a dot speed of 6.0º/s, dot density of 16.7 dots/(deg^2·s) and a circular aperture of 5.0º diameter), and the size and distance to center of the R-targets was adopted from Resulaj et al. (2009). Consequently, the possibility exists that a visual illusion might have emerged in previous experiments.

As reviewed in Section 2.1, until recently psychophysical and neuro-physiological studies on visual motion discrimination with random-dot motion stimuli were conducted as 2AFC tasks. With only two choice-

alternatives 180º apart, effects of the visual illusion would hardly be noticeable in the data. Accordingly, also in our data set the visual illusion became apparent only through a comparison of the performance between the 2-choice 90º- and 180º-cases and the 4-choice condition. Nonetheless, even in 2AFC tasks human subjects might have noticed the visual illusion, if they received feedback on the accuracy of their choice after each trial. Yet, to our knowledge no other study so far mentioned any illusion effects during the RDM task.

As for 2AFC tasks, also in the 3-alternative experiment of Niwa and Ditterich (2008) a visual illusion would barely become apparent. In their task, the three response targets were equally spaced 120º apart. Even if there had been a visual illusion in their RDM stimuli, the illusionary motion direction was not available as a possible choice.

The situation is different in the 2- and 4-choice experiments of Churchland et al. (2008). It would be very interesting to examine the monkeys' error responses in the 4-choice condition according to the relative proportions of 90º- and 180º-errors as in Fig. 6.4. Without the visual illusion, the 90º-errors should constitute 67% of the total errors for all coherence levels, as predicted by the attractor model without illusion bias ($\gamma = 1$).

One experimental characteristic that might have facilitated the visual illusion particularly in our case, concerns the distance of the response targets to the central fixation point. In most previous RDM studies, except for (Resulaj et al., 2009), subjects indicated their choice with a rather abrupt motor response, such as saccades or pressing a button. In order to observe change of mind, in our experiment and that of Resulaj et al. (2009) participants responded with a continuous hand movement. To allow for longer movement trajectories, R-targets were located further away from the central fixation point. If the R-targets in previous experiments were located close enough to the motion stimulus, they might have acted as reference points with respect to the dot motion. In this way, the illusionary impression that coherently moving dots seem stationary, while the randomly moving dots shift in the opposed direction, could have been impeded.

Resulaj et al. (2009) investigated changes of mind for two choice-alternatives. As stated above, there is little chance to find evidence of a hypothetical visual illusion in the 2-alternative data. One peculiarity in Resulaj et al.'s (2009) data might yet be indicative of the visual illusion: at 3.2% coherence all three tested subjects changed more from the correct to the wrong alternative than vice versa (see Fig. 3.2 for Subject S, where the effect was most pronounced). Note that for changes of mind based on random fluctuations in the motion evidence, erroneous changes are always less probable than correcting changes, at least for two choice-alternatives.

In sum, we cannot exclude the possibility that our particular experimental protocol and the combination of stimuli parameters, such as the more peripheral R-targets, might have provoked the visual illusion of reversed coherent motion. Nevertheless, our findings call attention to the general possibility of a visual illusion in the random-dot motion stimulus and care should be taken to avoid illusionary perceptions in future experiments.

Taken together, by means of the attractor model we showed that a hypothetical illusion of reversed coherent motion could have created the minor discrepancies that we encountered in the decision-behavior of our participants. By introducing a heterogeneous "illusion bias" in the model, we could account for the better performance in the 2-choice 90°-case compared to the 180°-case, and disproportional percentages of 180°-versus 90°- errors and changes of mind, without distorting any of the conclusions presented in the previous chapter. Finally, our findings imply that caution is required in future RDM experiments, especially in the context of decision-making beyond 2AFC tasks.

7 GENERAL DISCUSSION

In the last chapters we presented our investigations and findings on the neural computations that underlie sensorimotor decision-making with changes of mind and multiple choice-alternatives. Our universal approach was to apply biophysically realistic attractor models to behavioral and neurophysiological data of human and non-human primates performing different versions of the random-dot motion (RDM) task. Thereby, we encountered several common conditions and implications on the network states and dynamical regimes that enabled the models to reproduce the experimental observations. Here, we want to highlight these common findings and discuss general issues of our approach. In particular, we address the following questions: Are attractor states just a theoretical concept that provides explanations for observed neural dynamics, or is the brain really a dynamical system determined by neural attractors? Does the "biophysically realistic" attractor model include sufficient biological detail and is it accurate enough that its predictions can be taken seriously?

7.1 Are there attractors in the brain?

7.1.1 Findings in favor of attractor states

With this thesis we have confirmed the general applicability of attractor models to neural processes during decision-making, including changes of mind and multiple-choices. Apart from being congruent with the overall behavioral and neurophysiological findings, the attractor models could explain several experimental observations as direct consequence of their nonlinear attractor properties and the shape of the attractor landscape. In particular, for the 2-alternative model we found that the emergence of changes of mind was facilitated close to the "second bifurcation" between the region of bistability and the multi-stable regime (Fig. 3.6). There, the symmetric state becomes stable, in which both selective pools have high firing rates just below the decision threshold (Chapter 3). In Chapter 5 we confirmed this link between attractor states and changes of mind for multiple choices. Interestingly, the second bifurcation also played a role in accounting for Churchland et al.'s (2008) findings in the 2- and 4-alternative RDM task (Chapter 4). The longer reactions times and the shallower build-up activity that was observed experimentally in the 90°-control condition compared to the standard binary condition could be reproduced by the model with slightly stronger connections between neighboring selective pools. In that case, the

symmetric state for two neighboring pools is more attractive than between opposing pools, which delays the categorical decision in the 90º-case.

As we have discussed in Section 2.3, the attractor model also proved particularly appropriate to explain the nonlinear, time-dependent effects of motion pulses revealed in the experiments of Huk and Shadlen (2005), while DDM and LCA models could only explain the stronger influence of earlier motion pulses if explicit time-varying effects would be added to the decision threshold.

Similarly, in an RDM experiment in which the viewing duration was controlled by the experimenter, Kiani et al. (2008) found that neural activity and behavioral performance reached a plateau value for long stimulus times. Contrary to perfect integration, these findings can still be explained within the accumulation-to-bound framework if fixed decision thresholds are assumed to terminate the decision process as well in the interrogation paradigm. Here, the attractor model offers a complementary implicit explanation: when the system reaches a steady state, the decision attractor, the integration associated with neural build-up activity stops naturally at the plateau firing rate of the decision state.

7.1.2 A comprehensive account of diverse temporal dynamics

The above examples already indicate that attractor states are more than just an abstract theoretical concept for actual cortical processes. Moreover, unlike more phenomenological models such as the drift diffusion model, the biophysically realistic spiking-neuron attractor model aims at simulating realistic neural firing rates. In the context of the RDM task, simpler linear accumulation models can account for the build-up of activity observed in monkey LIP neurons during motion viewing prior to the motor choice (Ditterich, 2006b). Yet, LIP neurons, and decision-related neurons in general, display a variety of temporal dynamics during the time course of a RDM trial and also exhibit persistent activity in delayed decision tasks (e.g. Shadlen and Newsome, 2001). These seemingly peripheral aspects of LIP responses are not captured by phenomenological decision-making models (Mazurek et al., 2003), but provide important constraints on any model that attempts to give a comprehensive account of real neural computations (Wong and Huk, 2008).

As an exception among existing models of decision-making, attractor models can exhibit different modes of activity dependent on their current network state and dynamical regime. We have seen in Chapters 3 and 4, for the binary and multiple-choice case, that stable attractor states can be created and destroyed dependent on the inputs to the network (Fig. 3.6 and Fig. 4.5). If the attractor model is driven into the bistable regime, it

performs winner-take-all decision-making with slow transients to the decision-attractors. Moreover, the attractor model is capable of reproducing persistent activity with sufficiently strong recurrent connections, as in this case the model is in the multistable regime even without receiving any sensory inputs (see Fig. 2.8 and Fig. 4.A.2). Wong et al. (2007) further suggested that the symmetric state obtained for very high input currents to the selective populations, can explain LIP firing rates during the initial target phase before the onset of the RDM stimulus. In Chapter 4 we extended this idea to the case of four alternatives. Notably, the 2- and 4-choice attractor model accounted even quantitatively for the experimental finding that LIP activity during the target phase is lower if four R-targets are presented rather than just two.

Importantly, changes in the selective, sensory inputs are sufficient to shift the network from one dynamical regime to another and to induce transitions between stable states. Therefore, the same attractor network can subserve different functions during the time-course of a single trial, simply responding to the presented sensory stimuli (Wang, 2008).

What is more, this virtue of the attractor model is not restricted to the RDM paradigm. With a single line-attractor model Machens et al. (2005) could reproduce the graded delay activity and the subsequent categorical decision activity of PFC neurons during the different phases of a sequential vibrotactile discrimination task (see 2.1.2). The external signals in each task epoch automatically reconfigured the attractor landscape in such a way that it accounted first for working memory and subsequently for the decision computation.

Besides, working memory is in general strongly associated with self-sustained population activity patterns. In that context, stimulus-selective attractor states that arise through strong recurrent excitation form a leading candidate mechanism for the generation of mnemonic persistent activity (Amit, 1995; Wang, 2001; Brunel, 2003).

In sum, attractor models can comprehensively explain various temporal dynamics that have been observed in decision-related neurons by adopting different dynamical states. More importantly, transitions between these dynamical states can be induced in a way that is easily implementable in the brain, namely by increasing or decreasing the network inputs.

7.1.3 Adapting behavior through input

Along the same lines, in the attractor framework the brain could also make use of input changes to adjust decision-behavior, such as reaction times, accuracy, and changes of mind. Roxin and Ledberg (2008) showed for the reduced attractor model that in the proximity of the first bifurcation the speed-accuracy tradeoff could be controlled by changes in the

common input to both selective populations (Fig. 2.8 and Fig. 3.6). In Chapter 3 we verified this relation of faster reaction times for higher common inputs in the spiking-neuron model for the whole range of categorical decision-making, that is, the bistable regime. Moreover, with higher selective inputs the system approaches the second bifurcation where changes of mind are facilitated due to the symmetric state of elevated activity in both selective pools, which is consistent with our subsequent observation that faster reaction times correlate with more changes of mind (Chapter 5).

In the context of real cortical processes, the possibility to regulate the dynamical state and decision-behavior by means of the inputs to involved neural populations is particularly intriguing, as it offers a convenient explanation on how internal factors like attention, emotion, and reward expectancy might influence our decisions: Top-down inputs from higher-level brain regions could affect decision-making and, besides, also working memory in the same way as bottom-up sensory inputs (Deco and Rolls, 2005; Grabenhorst and Rolls, 2011).

7.1.4 Alternative approaches

Viewing the brain as a dynamical system with attractor states thus yields a comprehensive and physiologically implementable account for a wide range of cortical functions. Nevertheless, alternative approaches to working memory and decision-making have been suggested and implemented with spiking-neuron models. Mongillo et al. (2008), for instance, proposed a model of working memory based on short-term synaptic facilitation, where information can be maintained in memory in synaptic form without persistent activity. Yet, also in their model persistent activity can be obtained by making use of bistable attractor states. Hence, synaptic facilitation is not contrary, but rather complementary to the idea of attractors (see also Deco et al., 2010).

In the context of multiple-choice decisions, Beck et al. (2008) developed a continuous decision-making model based on probabilistic population codes, which we briefly discussed in Chapter 4. Although in their model the integration of evidence in layer LIP is performed linearly without attractor dynamics, attractors are employed for subsequent action selection.

7.1.5 Clinical implications

Whether attractor states play a crucial role in real neural processes, as suggested in this dissertation, could further be tested in pharmacological studies by impairing recurrent excitation or inhibition (see Section 3.3.5). On that note, it has been suggested recently that disorders in the stability

of neural attractor states underlie certain neural pathologies, like schizophrenia (Loh et al., 2007; Rolls et al., 2008b) and obsessive-compulsive disorder (Rolls et al., 2008a). Strikingly, pathological alterations in NMDA and GABA efficacies, which lead to shallower basins of attraction, i.e. to less stability, could collectively explain many seemingly inconsistent symptoms of schizophrenia from impaired working memory to hallucinations. Overstability of attractor states, on the other hand, could be related to repetitive actions and difficulties in switching to new actions, common to patients with obsessive-compulsive disorder. All in all, attractor states do seem to be essential for the proper functioning of cortical processes.

7.2 Is the attractor model realistic enough?

In the previous section we summarized a substantial number of indications that the brain indeed acts as a dynamical system with attractor states of neural activity. Even so, the biophysically realistic attractor networks we deployed in this thesis are still far from resembling a cortical region in all its details, with respect to both its basic neural units and network connections. This raises the question, whether and to what extent these models describe and predict neurophysiological processes adequately.

Aiming at a minimal representation of a cortical area, one could ask if spiking neurons are necessary at all in order to address our scientific objectives. Indeed, mean-field and firing-rate reductions of the attractor network with leaky integrate-and-fire (LIF) neurons have been proposed, as reviewed in Section 2.2.3 (Brunel and Wang, 2001; Wong and Wang, 2006; Roxin and Ledberg, 2008). These reductions can give valuable analytical insights. Yet, if the goal, as in our case, is to characterize and draw physiological conclusions about true dynamics of a finite-sized, noisy cortical network, there is no way around a spiking-neuron implementation (Deco et al., 2009).

7.2.1 Sparse connectivity and heterogeneous firing rates

At the network level, the presented spiking-neuron attractor models are subsampled reductions of real cortical networks, fully connected, and absolutely homogenous within one neural population. This has the computational advantage that fewer neurons have to be simulated. Yet, subsampling can also distort the network dynamics, for example by leading to artificial synchronization (Djurfeldt et al., 2008a). Just recently, the effect of sparse neural connections was tested by Rolls and Webb (Rolls and Webb, 2011), comparing the fully connected attractor model (Wang, 2002) to an otherwise identical model with diluted connectivity of

0.25 and 0.1 in the selective populations. They found that sparse networks are somewhat more stable, as dilution reduces finite size noise in the sparse networks, similar to increasing the number of neurons. The same authors also investigated how a graded, exponential firing-rate distribution within a neural population influences network properties (Webb et al., 2011). Interestingly, distributed heterogeneous firing rates have the opposite effect as sparseness: they increase stochasticity in the network. Nevertheless, in both cases the general dynamical properties and average population firing rates were very similar to the homogeneous, fully connected model and opposing effects on noise and stability in the network might even cancel out.

7.2.2 Physiological detail of neural units

In line with "Occam's razor", the attractor models we implemented take the approach to include as much physiological details as necessary to reproduce experimental findings, while at the same time keeping the model as simplistic as possible. Therefore, also the basic units, LIF neurons, are immense simplifications of real cortical pyramidal cells and interneurons to point-neurons without dendrites or ion channels. Generally, the desired degree of physiological detail in a network depends on the scientific questions posed. Here, we have been interested primarily in the dynamical properties and neural coding during decision-making, but not in any electrophysiological or pharmacological influences. In that case, threshold neuron models such as the LIF proved sufficient, as they give a very good account on the spike-timing of single neurons and neural assemblies (Sakai et al., 1999; Gerstner and Kistler, 2002; Gerstner and Naud, 2009).

In contrast to our "top-down" approach on modeling cortical processes, several projects have recently been launched with the objective to model entire cortical columns, or even whole brain areas, with as much neural detail as currently available from electrophysiological and imaging studies (Sporns et al., 2005; Markram, 2006; Izhikevich and Edelman, 2008; Deco et al., 2011). The hope of these projects is to capture functional properties by "reverse engineering" a full cortical column in all its known detail. Obtaining a mechanistic understanding of these models is almost as challenging as understanding the real brain. However, in contrast to the real tissue, an artificial cortical column could be investigated simply by changing parameters and observing the effects on the network's activity, i.e. running "virtual experiments". Yet, the computational load produced by these simulations still requires the latest massively parallel supercomputers.

Another problem of these "ultra-large-scale" models is that many important parameter values are still unknown. This in mind, it has to be

noted that a more detailed network is not necessarily more realistic: more parameters also introduce more uncertainty and more experimental data is required to constrain them. When empirical data is still missing, hypotheses from more abstract models can provide additional constraints (Djurfeldt et al., 2008a). As an example, the group of Prof. Anders Lansner takes a combined approach of bottom-up and top-down modeling. They investigate memory attractor states in increasingly detailed, full-scale models of large patches of cortical layers II/III, which are comprised of multi-compartment Hodgkin-Huxley neurons (Djurfeldt et al., 2008b; Lansner, 2009). These model neurons explicitly simulate spike dynamics based on different ionic currents. Using this type of neurons thus allows testing pharmacological, chemical and temperature effects on memory function.

Taken together, the type of "biophysically realistic" attractor models that we used in this thesis are minimalistic models, including physiological details with a functional purpose. They successfully account for neurophysiological processes, but are still sufficiently abstract to allow tracing the dynamics and understanding the general mechanisms at work. On that note, the visual illusion in the RDM task (Chapter 6), which we revealed with the help of the attractor model, is a nice example of the models' explanatory power.

8 CONCLUSION AND OUTLOOK

In this thesis we have combined several studies with the common purpose of investigating neuronal computations behind perceptual decision-making, beyond the usual restriction to two-alternative forced-choice tasks. Our special focus was on sensorimotor decisions with multiple choice-alternatives and the possibility for changes of mind. As an overall conclusion of our work, we confirmed that established decision-making models and, in particular, biophysically realistic attractor models, can be extended to account for these more complex aspects of decision-making.

Furthermore, we were able to show that the experimental findings from random-dot motion discrimination tasks with changes of mind and multiple choices provide additional constraints to existing theoretical models of decision-making and are thus informative about the neural decision process in general. Especially the way how the speed-accuracy tradeoff and changes of mind are interrelated and depend on the decision threshold, shared sensory inputs, and the number of choice-alternatives, might further be exploited to distinguish linear from nonlinear modeling approaches in future experiments.

The 2- and 4-alternative RDM experiment of Churchland et al. (2008), for instance, indicated that the decision threshold is independent of the number of choice-alternatives. Longer reaction times in the case of four alternatives seem to be caused by a lower initial level of neural activity just prior to the gradual activity build-up associated with evidence accumulation. Previously, theoretical models reproduced increasingly longer reaction times for more choice-alternatives with different decision thresholds (Usher and McClelland, 2001; Usher et al., 2002). The discrete 2- and 4-choice attractor model we presented in Chapter 4, however, accounts for the conditional differences in activity before motion onset. For the condition that the regime of categorical decision-making coincides for the different numbers of alternatives, global inhibition in the network proved sufficient to reproduce the neurophysiological findings in the 2- and 4-alternative condition. The model further suggests that a larger population size of the selective pools, which represent the choice-alternatives, is advantageous in order to obtain the postulated overlap of decision regimes without further top-down mechanisms. What is more relevant, resulting from its particular nonlinear attractor properties, our discrete 2- and 4-choice attractor network is so far the only proposed model that captures the experimental observation of the 90°-control case, using the same decision threshold as for the standard 2- and 4-choice conditions.

Besides, up to now, the neural decision threshold in LIP neurons during the RDM task has been estimated by the neural activity with least variability across trials just prior to the saccadic motor response. In Chapter 3 we have proposed that also the mean neural activity from changes-of-mind trials might be used to assess the decision threshold. More precisely, the biophysically realistic attractor model predicts a switch in neural activity during changes of mind, which could be tested experimentally and might be indicative of the neural decision threshold (Fig. 3.4).

Furthermore, the decision threshold is also involved in adjusting the speed-accuracy tradeoff. Interestingly, a relation between changes of mind and the speed-accuracy tradeoff was found experimentally and for the attractor model for binary and multiple choices. In Chapters 3 and 5 we demonstrated that speed pressure leads to more changes of mind and reaction times are negatively correlated with the number of changes. While in linear, phenomenological models, such as the drift-diffusion model, only the decision threshold regulates the speed-accuracy tradeoff, in nonlinear attractor models the speed-accuracy tradeoff can also be adjusted by the synaptic input to the selective populations (see Chapter 3 and Roxin and Ledberg, 2008). In the bistable regime of categorical decision-making, higher common inputs lead to faster reaction times. Higher inputs also shift the dynamical working point of the system closer towards the second bifurcation at the border to the multistable regime where the symmetric state is stable (Fig. 3.6). Notably, we found that it is just close to this second bifurcation that changes of mind arise naturally in the biophysically realistic attractor model (Chapter 3). Correspondingly, we showed that in the attractor model the number of changes can be controlled in the same two ways as the speed-accuracy tradeoff: more changes of mind can be obtained through a lower decision threshold and higher common inputs.

Given these dependencies between the network input and decision behavior, we suggest that the brain could take advantage of the whole range of inputs, which enable categorical decision-making, in order to adjust decision speed, accuracy, and flexibility in the form of changes of mind.

As we have noted in our model review in Section 2.3, one of the great future challenges in the field of perceptual decision-making is to design novel experiments in order to distinguish between competing modeling approaches. In this dissertation, we have shown that extensions of classic 2AFC paradigms, in experiment and theory, can yield substantial new insights into the neural implementation of decision-making. Moreover, our findings allowed us to propose future experiments with the aim to further evaluate the validity of the attractor model and to distinguish between the linear drift-diffusion model and the nonlinear attractor model.

In particular, in Chapter 3 we have proposed a RDM experiment with independent coherent motion in two opposed directions to distinguish between the two models. In this way, the total amount of sensory evidence for both decision alternatives could be altered, while keeping the bias fixed. It is important to note that, contrary to the attractor model, the decision behavior of the linear drift-diffusion model would not be affected if the evidence in favor of both alternatives was increased by the same amount. Due to differences in the input variance, it might still be difficult to exclude one of the models based on behavioral data alone. Yet, the two models make clearly distinguishable predictions on the neuro-physiological level (Fig. 3.7 and Fig. 3.8).

Another promising future endeavor would certainly be to extend the investigations on the relation between changes of mind and the number of choice-alternatives. To this point, we have demonstrated that choice corrections, and, possibly, changes of mind in general, become less likely with more choice alternatives, in accordance with the attractor model (Chapter 5). More significant relations, based on larger conditional differences, might be gained if participants with extensive training were tested on the same, or a similar, experiment as reported in Chapter 5. Interactions between experimental conditions, changes of mind, and the number of choice alternatives could further confirm and challenge current modeling approaches.

To conclude, the neurophysiological and behavioral findings, theoretical/computational explanations and predictions, and also the further experimental evaluations that we have presented in this thesis complement and extend our current understanding of perceptual decision-making. Indeed, we believe that this study reflects the timely necessity to turn from the most basic and phenomenological models of decision-making to more realistic models, with an increasing degree of neurophysiological detail, in order to advance our understanding of the processes underlying decision-making.

A APPENDIX

A.1 Theoretical Framework

A.1.1 Detailed mathematical description of general model characteristics

The attractor model with spiking neurons and biophysically realistic synaptic conductances and receptors was first introduced by Brunel and Wang (2001). It was originally intended to model object working memory and persistent activity, but also successfully simulated behavioral and neurophysiological data obtained during a binary RDM task (Wang, 2002) and perceptual decision-making in general (Deco and Rolls, 2006; Wang, 2008).

a) Network

Single neurons are modeled as leaky integrate-and-fire neurons (LIF) with conductance-based synaptic responses (see below). They are connected by three types of receptors that mediate the synaptic currents flowing into them: AMPA, NMDA glutamate, and $GABA_A$ receptors, which are described by realistic synaptic kinetics (Eq. A.3-9).

The network is fully connected and divided into a population of excitatory neurons (80%) and inhibitory neurons (20%) (Abeles, 1991). Some of the excitatory neurons are thought to respond selectively to one of the possible directions of coherent motion used in the experiment and, hence, reflect the possible alternatives. Thus, the excitatory neurons are subdivided into several "selective" neural populations (pools) and one further pool of nonselective neurons. The nonselective pool emulates the activity in the surrounding brain areas. Each selective pool contains $f = N_E$ neurons. The fraction f is called the "coding level" of the selective pools. Neuronal pools generally are defined as groups of neurons sharing the same inputs and connectivities. The inhibitory pool is homogeneous and regulates the overall activity by implementing competition in the network.

b) Synaptic weights

The synaptic efficacies are assumed to be already formed and, therefore, kept fixed during the simulation. They are consistent with a Hebbian rule: the synapse between two cells is strong if their activity was correlated in the past, low if it was anticorrelated. The baseline connection weight between uncorrelated excitatory populations is set to 1. Cells within one selective pool have stronger recurrent connection weights ($\omega_+ > 1$), as their activity is thought to be more correlated. Cells between

143

selective pools and from the nonselective to selective pools with rather anticorrelated activity have weaker connection weights (ω_- < 1). The total incoming connections to each selective neuron are normalized to 1 by adjusting ω_- relative to ω_+ and the coding level f (A.1.3 and A.1.5). In this way, the overall excitatory recurrent synaptic drive remains constant in the spontaneous state, despite changes in the network parameters (Brunel and Wang, 2001). Inhibitory connections are denoted by a weight ω_I. The specific connection weights used in the different model versions are summarized in Table A.1-6. For simplicity, we used instantaneous synapses without delays.

c) Spiking dynamics

Neurons. LIF neurons (Tuckwell, 1988) do not explicitly model action potentials, or "spikes", but give a realistic account of the sub-threshold membrane potential, characterized by the following equation:

$$C_{\mathrm{m}}\frac{dV(t)}{dt} = -g_{\mathrm{m}}(V(t) - V_{\mathrm{L}}) - I_{\mathrm{syn}}(t)$$

(A.1)

with resting potential V_L, a membrane capacitance C_m and a membrane leak conductance g_m. I_{syn} is the total synaptic current flowing into the cell.

When the membrane potential of a LIF neuron reaches the firing threshold V_{th}, a spike is registered and the membrane potential is reset to V_{reset}, and clamped there during an absolute refractory period of τ_{ref}.

Synapses. Recurrent excitatory post-synaptic currents (EPSCs) are mediated by fast AMPA and slow NMDA glutamate receptors, inhibitory post-synaptic currents (IPSCs) by GABA$_A$ receptors. External inputs are assumed to arrive only via fast AMPA receptors. The total synaptic current is thus given by the sum:

$$I_{\mathrm{syn}}(t) = I_{\mathrm{AMPA,rec}}(t) + I_{\mathrm{NMDA,rec}}(t) + I_{\mathrm{GABA}}(t) + I_{\mathrm{AMPA,ext}}(t)$$

(A.2)

Individual synaptic currents are defined by:

$$I_{\mathrm{AMPA,rec}}(t) = g_{\mathrm{AMPA,rec}}(V(t) - V_{\mathrm{E}})\sum_{j=1}^{N_{\mathrm{E}}}\omega_j s_j^{\mathrm{AMPA,rec}}(t)$$

(A.3)

$$I_{\mathrm{NMDA,rec}}(t) = \frac{g_{\mathrm{NMDA}}(V(t) - V_{\mathrm{E}})}{1 + [\mathrm{Mg}^{2+}]exp(-0.062V(t))/3.57} \times \sum_{j=1}^{N_{\mathrm{E}}}\omega_j s_j^{\mathrm{NMDA}}(t)$$

(A.4)

$$I_{\mathrm{GABA}}(t) = g_{\mathrm{GABA}}(V(t) - V_{\mathrm{I}})\sum_{j=1}^{N_{\mathrm{I}}}\omega_j s_j^{\mathrm{GABA}}(t)$$

(A.5)

$$I_{\mathrm{AMPA,ext}}(t) = g_{\mathrm{AMPA,ext}}(V(t) - V_{\mathrm{E}})\sum_{j=1}^{N_{\mathrm{ext}}} s_j^{\mathrm{AMPA,ext}}(t)$$

(A.6)

with excitatory and inhibitory reversal potentials V_E, V_I, and synaptic weights ω_j. Sums add over all synapses formed by presynaptic neuron j.

The NMDA current is potential dependent, controlled by the extracellular concentration of magnesium $[Mg^{2+}]$ = 1 mM (Jahr and Stevens, 1990). The fraction of open channels or gating variable s_j for AMPA and GABA$_A$ receptor mediated currents is described by

$$\frac{ds_j(t)}{dt} = -\frac{s_j(t)}{\tau_{\mathrm{decay}}} + \sum_k \delta(t - t_j^k)$$
(A.7)

where the rise time constants have been neglected because they are smaller than 1 ms. This is not possible for NMDA mediated currents with a rise time constant of $\tau_{NMDA,rise}$ = 2 ms. There s_j is determined by:

$$\frac{ds_j^{\mathrm{NMDA}}(t)}{dt} = -\frac{s_j^{\mathrm{NMDA}}(t)}{\tau_{\mathrm{NMDA,decay}}} + \alpha x_j(t)(1 - s_j^{\mathrm{NMDA}}(t))$$
(A.8)

$$\frac{dx_j(t)}{dt} = -\frac{x_j(t)}{\tau_{\mathrm{NMDA,rise}}} + \sum_k \delta(t - t_j^k)$$
(A.9)

with α = 0.5 kHz. The sums over k represent a sum over the spike train t_k with δ-Peaks $\delta(t)$ emitted by presynaptic neuron j at time t_j^k. The values for the decay time constants are τ_{AMPA} = 2 ms for AMPA synapses, $\tau_{NMDA,decay}$ = 100 ms for NMDA synapses (Hestrin et al., 1990; Spruston et al., 1995) and τ_{GABA} = 10 ms for GABA synapses (Salin and Prince, 1996; Xiang et al., 1998).

The neuronal and synaptic capacities and time constants of our model are taken from the original model of persistent activity (Brunel and Wang, 2001). The conductances there were calibrated in order to obtain a physiological spontaneous firing rate of 3 Hz for excitatory neurons and 9 Hz for inhibitory neurons in the unstructured network. As we used different total numbers of neurons in each model version, the recurrent conductances used by Brunel and Wang (2001) had to be scaled by a factor of 1000/#neurons to keep the mean recurrent input constant. In addition, in the model of Brunel and Wang (2001), recurrent excitation is largely mediated by NMDA receptors, taking advantage of their slower synaptic dynamics to stabilize the sustained activity state. In our model versions we could accomplish better approximations to the experimental data by slightly increasing the amount of AMPA relative to NMDA. Because in the original model the effective NMDA/AMPA ratio near firing threshold is 10 in terms of charge entry, a decrease in g_{NMDA} has to be compensated by a tenfold increase in g_{AMPA} in order to preserve the spontaneous firing rate:

$$g_{\mathrm{NMDA}} = g_{\mathrm{NMDA}}(1 - \delta), \quad g_{\mathrm{AMPA}} = g_{\mathrm{AMPA}}(1 + 10\delta).$$
(A.10)

d) Network inputs

External inputs are generally modeled as uncorrelated Poisson spike trains applied independently to the individual neurons in the network. All neurons receive a background input of v_{ext} = 2.4 kHz, equivalent to 800 excitatory connections from external neurons firing at 3 Hz. It simulates spontaneous noisy background activity from outside the local network and sensory information processed by other brain areas. In all our simulations, the network is given 500 ms, where it receives only the background inputs, to adjust from the initially set firing rates to the true spontaneous state.

Subsequently, on top of the background activity an external target and motion input are applied to the selective neural populations only. They correspond to the sensory stimuli during the RDM experiment: the visually shown R-targets and the random-dot motion respectively. The specific input functions are described in detail for each model version in Sections A.1.3,5 and 6.

e) Decision threshold and non-decision time

According to recent experimental findings (Roitman and Shadlen, 2002; Churchland et al., 2008), we assumed fixed decision thresholds independent of motion coherence.

Reaction times were calculated as the time of threshold crossing plus a non-decision time t_{ND}, which consists of the latency the motion signal needs to arrive in area LIP (or PCC in general), and the duration of movement initiation and execution. We generally assumed a motion signal latency of 200 ms (Roitman and Shadlen, 2002; Churchland et al., 2008). To simulate the time for hand movements in Chapter 3 and 5 we used a value of 180 ms (Snyder et al., 1997; Cui and Andersen, 2007), for saccade initiation and execution in Chapter 4 we used 80 ms according to (Roitman and Shadlen, 2002; Churchland et al., 2008). In the simulations that included changes of mind, the non-decision time also set the time limit for the changing.

A.1.2 Mean-field approximation

In the mean-field approximation the number of integration variables is reduced to one for each neural population (Brunel and Wang, 2001). Solving the mean-field equations provides the fixed points of the population firing rates, i.e. the stationary states of the populations after the period of dynamical transients. As this can be done much more quickly than integrating the full spiking model, scanning the parameter space in order to find a parameter set matching the experimental findings becomes feasible.

In the mean-field formulation the potential of a neuron is calculated according to:

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

(A.11)

where $V(t)$ is the membrane potential, x labels the populations, τ_x is the effective membrane time constant, μ_x is the mean value the membrane potential would have in the absence of spiking and fluctuations, σ_x measures the magnitude of the fluctuations and η is a Gaussian process with exponentially decaying correlation function and time constant τ_{AMPA}. The quantities μ_x and σ_x^2 are given by:

$$\mu_x = \frac{(T_{\text{ext}}\nu_{\text{ext}} + T_{\text{AMPA}}n_x^{\text{AMPA}} + \rho_1 n_x^{\text{NMDA}})V_{\text{E}} + \rho_2 n_x^{\text{NMDA}}\langle V \rangle + T_{\text{I}}n_x^{\text{GABA}}V_{\text{I}} + V_{\text{L}}}{S_x}$$

(A.12)

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

(A.13)

where ν_{ext} is the external incoming spiking rate, ν_I is the spiking rate of the inhibitory population, $\tau_m = C_m/g_m$ with the values for the excitatory or inhibitory neurons depending of the population considered. The other quantities are given by:

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

(A.14)

$$\tau_x = \frac{C_{\text{m}}}{g_{\text{m}}S_x}$$

(A.15)

$$n_x^{\text{AMPA}} = \sum_{j=1}^{p} f_j w_{jx}^{\text{AMPA}} \nu_j$$

(A.16)

$$n_x^{\text{NMDA}} = \sum_{j=1}^{p} f_j w_{jx}^{\text{NMDA}} \psi(\nu_j)$$

(A.17)

$$n_x^{\text{GABA}} = \sum_{j=1}^{p} f_j w_{jx}^{\text{GABA}} \nu_j$$

(A.18)

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

(A.19)

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

(A.20)

$$\tau_{\text{NMDA}} = \alpha\tau_{\text{NMDA,rise}}\tau_{\text{NMDA,decay}}$$

(A.21)

$$T_{\text{ext}} = \frac{g_{\text{AMPA,ext}}\tau_{\text{AMPA}}}{g_m}$$

(A.22)

$$T_{\text{AMPA}} = \frac{g_{\text{AMPA,rec}}N_{\text{E}}\tau_{\text{AMPA}}}{g_m}$$

(A.23)

$$\rho_1 = \frac{g_{\text{NMDA}} N_{\text{E}}}{g_m J} \tag{A.24}$$

$$\rho_2 = \beta \frac{g_{\text{NMDA}} N_{\text{E}} (\langle V_x \rangle - V_{\text{E}})(J-1)}{g_{\text{m}} J^2} \tag{A.25}$$

$$J = 1 + \gamma \exp(-\beta \langle V_x \rangle) \tag{A.26}$$

$$T_{\text{I}} = \frac{g_{\text{GABA}} N_{\text{I}} \tau_{\text{GABA}}}{g_{\text{m}}} \tag{A.27}$$

$$\langle V_x \rangle = \mu_x - (V_{\text{thr}} - V_{\text{reset}}) \nu_x \tau_x \ , \tag{A.28}$$

where p is the number of excitatory populations, f_x is the fraction of neurons in the excitatory population x, $\omega_{j,x}$ the weight of the connections from population x to population j, ν_x is the spiking rate of the excitatory population x, $\gamma = [\text{Mg}^{2+}]/(3.57 \text{ mM})$, $\beta = 0.062 \text{ mV}^{-1}$ and the average membrane potential $\langle V_x \rangle$ has a value between -55mV and -50 mV.

The mean field approximation finally yields a set of n nonlinear equations describing the average firing rates of the different populations in the network as a function of the defined quantities μ_x and σ_x:

$$\nu_x = \phi(\mu_x, \sigma_x), \quad x = 1, ..., n, \tag{A.29}$$

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

$$\phi(\mu_x, \sigma_x) = \left(\tau_{rp} + \tau_x \int_{\beta(\mu_x, \sigma_x)}^{\alpha(\mu_x, \sigma_x)} du \sqrt{\pi} \exp(u^2)[1 + \text{erf}(u)] \right)^{-1} \tag{A.30}$$

$$\alpha(\mu_x, \sigma_x) = \frac{(V_{\text{thr}} - \mu_x)}{\sigma_x} \left(1 + 0.5 \frac{\tau_{AMPA}}{\tau_x} \right) + 1.03 \sqrt{\frac{\tau_{AMPA}}{\tau_x}} - 0.5 \frac{\tau_{AMPA}}{\tau_x} \tag{A.31}$$

$$\beta(\mu_x, \sigma_x) = \frac{(V_{reset} - \mu_x)}{\sigma_x} , \tag{A.32}$$

with $erf(u)$ the error function and τ_{rp} the refractory period which is considered to be 2 ms for excitatory neurons and 1 ms for inhibitory neurons. To solve the equations defined by Eq. A.29 for all x, we numerically integrate Eq. A.28 and the differential equation below, whose fixed-point solutions correspond to solutions to Eq. A.29:

$$\tau_x \frac{d\nu_x}{dt} = -\nu_x + \phi(\mu_x, \sigma_x) \tag{A.33}$$

To find the possible fixed points that coexist for a given parameter set, Eq. A.33 has to be integrated for different initial conditions of population firing rates over a range of external inputs. Generally, the firing rates obtained by the mean-field approximation would be exact if

the number of neurons was infinitely large and the unitary postsynaptic potentials elicited by presynaptic spikes were infinitesimally small.

A.1.3 Model specifications for binary changes of mind

The attractor network presented in Chapter 3 has the same general structure as the original study on binary decision-making (Wang, 2002). The network kinetics are summarized in Table A.1. With two selective pools in the network and a coding level of $f = 0.2$, ω_- is calculated as $\omega_- = (1 - f\omega_+)/(1 - f)$. To account for changes of mind in the binary RDM task (Resulaj et al., 2009), we adapted the weight parameters and inputs within biologically plausible boundaries (see below). All default simulation parameters are listed in Table A.2. Compared to the original studies (Brunel and Wang, 2001; Wang, 2002), the AMPA/NMDA ratio was increased according to Eq. A.10 with $\delta = 0.08$.

a) Simulation and analysis details

1,000 trials of 3,500 ms with different random seeds were run for each parameter set and motion coherence.

Decision thresholds. In the main simulations of Chapter 3, a (first) decision was reached when one selective pool crossed a threshold of 44 Hz and surpassed the other by at least 10 Hz. The same conditions applied for a change of mind. To confirm the conclusions from the mean-field approximation, additional simulations were run with different target inputs after motion input onset (from 1,500 ms on 25 Hz and 125 Hz instead of 85 Hz), and also for higher and lower inhibitory weights ($\omega_I = 1.425$ and $\omega_I = 0.825$ instead of 1.125 as in the standard simulations). The respective threshold values were: 30 Hz for 25 Hz target input, 50 Hz for 125 Hz target input, 38 Hz for the simulations with $\omega_I = 1.425$ and 50 Hz for $\omega_I = 0.825$. All threshold values used were determined within 1 Hz accuracy in order to match the experimental reaction times and percentage of correct choices of (Resulaj et al., 2009). A threshold alteration of ± 1 Hz roughly corresponds to a $\pm 3\%$ variation in reaction time and about $\mp 10\%$ in the frequency of changes). For the simulations shown in Fig. 3.7 of Section 3.3.6, the standard threshold parameters were used (44 Hz with 10 Hz difference). The additional condition of a minimal difference of 10 Hz between the firing rates of the two selective populations avoids occasional joint crossings to count as decisions or changes (Fig. 3.A.2).

Fluctuation analysis. The input firing rates for the fluctuation analysis of the external Poisson inputs (Section 3.3.3) were determined by filtering the external input spikes in the same way as the output spikes when calculating population firing rates, namely by averaging over a 50 ms time window, shifted with a time step of 5 ms.

Table A.1 Binary attractor model for changes of mind

A	Model Summary
Populations	Four
Topology	One module
Connectivity	Full connectivity, no synaptic delay
Neuron model	Leaky integrate-and-fire neurons, fixed voltage threshold, fixed absolute refractory periods
Channel models	-
Synapse models	Conductance-based synapses, AMPA and GABA$_A$ receptors (instantaneous rise, exponential decay), voltage-dependent NMDA receptors (exponential rise and decay)
Plasticity	-
Input	Independent fixed-rate poisson spike trains to all neurons
Measurements	Spike activity

B	Populations				
Total number of neurons	$N = 1000$		Excitatory neurons	$N_E = 0.8 \cdot N$	
			Inhibitory neurons	$N_I = 0.2 \cdot N$	

Name	**Size**	**Name**	**Size**
Selective pool 1 (right)	$N_{S1} = f \cdot N_E$	Nonselective	$(1 - 2f) \cdot N_E$
Selective pool 2 (left)	$N_{S2} = f \cdot N_E$	Inhibitory	$0.2 \cdot N$

C		Connectivity			
Source	**Target**	**Weight**	**Source**	**Target**	**Weight**
inhibitory $\mapsto$	all	$\omega_I = 1.125$	nonselective $\mapsto$	selective	$\omega_- = 0.8725$
excitatory $\mapsto$	inhibitory	$\omega = 1$	selective i $\mapsto$	selective j	$\omega_- = 0.8725$
excitatory $\mapsto$	nonselective	$\omega = 1$	selective i $\mapsto$	selective i	$\omega_+ = 1.51$

D	Neuron and Synapse Model
Type	Leaky integrate-and-fire neurons, conductance-based synapses
Subthreshold dynamics	$C_{\mathrm{m}} \dot{V}(t) = -g_{\mathrm{m}}(V(t) - V_{\mathrm{L}}) - I_{\mathrm{syn}}(t)$
Synaptic currents	$I_{\mathrm{syn}}(t) = I_{\mathrm{AMPA,rec}}(t) + I_{\mathrm{NMDA,rec}}(t) + I_{\mathrm{GABA}}(t) + I_{\mathrm{AMPA,ext}}(t)$
	$I_{\mathrm{AMPA,ext}}(t) = g_{\mathrm{AMPA,ext}}(V(t) - V_{\mathrm{E}}) \sum_{j=1}^{N_{\mathrm{ext}}} s_j^{\mathrm{AMPA,ext}}(t)$
	$I_{\mathrm{AMPA,rec}}(t) = g_{\mathrm{AMPA,rec}}(V(t) - V_{\mathrm{E}}) \sum_{j=1}^{N_{\mathrm{E}}} \omega_j s_j^{\mathrm{AMPA,rec}}(t)$
	$I_{\mathrm{NMDA,rec}}(t) = \frac{g_{\mathrm{NMDA}}(V(t) - V_{\mathrm{E}})}{1 + [\mathrm{Mg}^{2+}] exp(-0.062 V(t))/3.57} \times \sum_{j=1}^{N_{\mathrm{E}}} \omega_j s_j^{\mathrm{NMDA}}(t)$
	$I_{\mathrm{GABA}}(t) = g_{\mathrm{GABA}}(V(t) - V_{\mathrm{I}}) \sum_{j=1}^{N_{\mathrm{I}}} \omega_j s_j^{\mathrm{GABA}}(t)$
Fraction of open channels	$\frac{ds_j^{\mathrm{X}}(t)}{dt} = -\frac{s_j^{\mathrm{X}}(t)}{\tau_{\mathrm{X,decay}}} + \sum_k \delta(t - t_j^k), \quad \mathrm{X = AMPA, GABA}$
	$\frac{ds_j^{\mathrm{NMDA}}(t)}{dt} = -\frac{s_j^{\mathrm{NMDA}}(t)}{\tau_{\mathrm{NMDA,decay}}} + \alpha x_j(t)(1 - s_j^{\mathrm{NMDA}}(t))$
	$\frac{dx_j(t)}{dt} = -\frac{x_j(t)}{\tau_{\mathrm{NMDA,rise}}} + \sum_k \delta(t - t_j^k)$
Spiking	if $V(t) \geq V_{\mathrm{th}} \wedge t > t^\star + \tau_{\mathrm{ref}}$ 1. set $t^\star = t$ 2. emit spike at time $t^\star$ 3. set $V(t) = V_{\mathrm{reset}}$

E	Input	
Type	**Description**	
Poisson generator	Fixed rate $\nu = \nu_{ext} + \nu_{1,2}$, with ν_{ext} = 2.4 kHz, $\nu_{1,2}$ = sel. inputs, one generator per neuron	

F	Measurements
Spike activity: firing-rates were calculated using the spike count in a 50 ms time window shifted by 5 ms steps and dividing by the number of neurons in the population and by the window size.	

The "variation from mean input difference" (Fig. 3.5) was calculated by subtracting the mean input rate across trials from each selective population. The remaining input difference between the selective populations in each trial was then signed with respect to the first pool that crossed the decision threshold.

Fixed points in the spiking network. To obtain the stable states of the standard spiking-neuron model in comparison to the mean-field analysis (Fig. 3.6A, blue crosses), we simulated 100 trials each, without target inputs, but for constant symmetric inputs to the selective populations, ranging from 0 to 200 Hz in steps of 10 Hz for 3,500 ms. The stable fixed points of the decision state were found by averaging the last 500 ms of all trials in which the decision attractor was reached. For (very) low and high inputs, in some (most) of the trials the symmetric spontaneous or double-up state was stable and no decision was formed. The mean firing rate from 1,000 to 2,000 ms of these trials determined the fixed point of the respective symmetric state.

b) Mean-field analysis

As the mean-field approximation provides the fixed points of the attractor network, it enabled us to analyze the location of the network's working point with respect to the two crucial bifurcations that contain the range of categorical decision-making. Thereby, it also helped to find a parameter set to replicate the behavioral data on changes of mind (Resulaj et al., 2009).

Stable fixed points were found by terminating integration when the firing rates did not differ by more than 10^{-8} from the mean over the last 40 ms. Unstable fixed points were determined by the boundary of the basins of attraction between two stable states, searched by iterating the initial values between two stable branches to find the change of dynamic flow towards one or the other stable state. To find all possible fixed points that coexist for a given parameter set, we integrated Eq. A.33 with different initial conditions of population firing rates over a range of external inputs from 0 to 200 Hz in steps of 1.0 Hz.

A.1.4 Diffusion model for binary changes of mind

The results shown in Fig. 3.8 (Section 3.3.6) were obtained by numerically integrating[19] a diffusion model with an added second threshold and time-out for changing as described in (Resulaj et al., 2009). For the drift and boundary parameters, we used the average fitted values of Subject S from Resulaj et al. (2009): a drift rate $\mu = coh \cdot k$, with k = 0.3, a first decision bound $B = 13.2$, $t_{ND} = 324$ ms and $B\Delta = 23.3$, without any

---

[19] We used the Euler method with a step size *dt* of 1 ms.

Table A.2 Parameter set of the binary attractor model for changes of mind

Parameter	Value	Parameter	Value
N_E	800	N_I	200
V_E	0 mV	V_I	-70 mV
V_L	-70 mV	N_{ext}	800
V_{thr}	-50 mV	ν_{ext}	2.4 kHz
V_{reset}	-55 mV	f	0.20
ω_+	1.51	ω_I	1.125
ω_-	0.8725	α	$0.5\ \text{ms}^{-1}$
C_m (excitatory)	0.5 nF	C_m (inhibitory)	0.2 nF
g_m (excitatory)	25 nS	g_m (inhibitory)	20 nS
τ_{ref} (excitatory)	2 ms	τ_{ref} (inhibitory)	1 ms
$g_{AMPA,ext}$ (excitatory)	2.08 nS	$g_{AMPA,ext}$ (inhibitory)	1.62 nS
$g_{AMPA,rec}$ (excitatory)	0.1872 nS	$g_{AMPA,rec}$ (inhibitory)	0.1458 nS
g_{GABA} (excitatory)	1.25 nS	g_{GABA} (inhibitory)	0.973 nS
g_{NMDA} (excitatory)	0.30084 nS	g_{NMDA} (inhibitory)	0.23736 nS
$\tau_{NMDA,decay}$	100 ms	$\tau_{NMDA,rise}$	2 ms
τ_{AMPA}	2 ms	τ_{GABA}	10 ms

bias in starting point or drift ($\mu_0 = 0$, $y_0 = 0$). The increments of evidence were obtained from normal distributions with several variance levels. To obtain the predictions on alterations in input variance, we simulated 10,000 trials for each of the six coherence levels, with input variances of 0.7, 1.0 and 1.3, respectively, at time steps of 1 ms.

A.1.5 Multiple choice model for primate data

In Chapter 4, we extended the binary attractor model of decision-making to up-to-four possible choice-alternatives (Fig. 4.1B). Instead of two, there are now four selective populations, encoding the four possible directions of coherent motion. In the standard simulations fitting the experimental data of Churchland et al. (2008) the coding level is $f = 0.2$. Moreover, in order to model the spatial distribution of the R-targets in the experiments, we introduced a spatial connectivity component ω_T, which was added to the recurrent connection-weights between "neighboring" selective populations (Fig. 4.1C).

Consequently, ω_- is calculated as $\omega_- = 1 - f(2\omega_T + \omega_+ - 1)/(1 - f)$ to normalize the overall excitatory recurrent weights towards each neuron. The neural and synaptic equations for the 4-alternative model are identical to Eq. A.1-9. All other model details are summarized in Table A.3. All default simulation parameters are listed in Table A.4. Compared to the original studies (Brunel and Wang, 2001; Wang, 2002), the AMPA/NMDA ratio was increased according to Eq. A.10 with $\delta = 0.1$.

a) Simulation details

Each trial in the network was run for a total of 4,000 ms. For each parameter set a block of 1,000 trials with different random seeds was simulated. A decision was reached when one selective pool crossed a threshold of 50 Hz and surpassed the other selective pools by at least 5 Hz. For the 4-alternative condition in some cases, for small coherence levels, no decision was reached within the 4,000-ms simulation, as all selective pools stayed in the spontaneous state of low firing rate (at most 20 of 1,000 trials). These trials were discarded as failed trials and excluded from the average. In the 90°-case on the other hand, some trials had to be excluded, because the two neighboring selective pools stayed in a symmetric double state of enhanced activity and no decision was made (at most 54 of 1,000 trials, for low coherence levels).

b) Mean-field approximation

In the multiple-choice model, the mean-field approximation allowed us to determine the conditions where the competition regimes for the 2- and 4-choice cases overlap. To find all of the possible fixed points that coexist for a given parameter set, we integrated Eq. A.33 with different initial conditions of population firing rates. We used four initial conditions spanning the possible firing rates at different temporal stages of the spiking simulation and the different experimental conditions:

1. all selective pools with an initial firing rate of 0 Hz,
2. one selective pool 120 Hz, the other 3 pools 0 Hz,
3. two opposing selective pools 30 Hz, the other 20 Hz,
4. all selective pools with 30 Hz.

For our modeling purpose, decision-making between 2- and 4-alternatives, we intended to find a region of multistability with competition, so that always just one selective pool would terminate in an up-state of high firing rate. For specific sets of model parameters, the stable fixed points were calculated over a range of external inputs from 0 to 100 Hz in steps of 0.5 Hz for the 2- and 4-choice condition (Fig. 4.5A). The range of external inputs where for all initial conditions and both experimental paradigms a decision is reached, i.e., one and just one pool is in an up-state, was termed ''range of decision-making.'' Keeping the other parameters fixed, the value of the neighboring connectivity ω_T with the optimal, i.e., broadest, range of decision-making was determined by performing the fixed-point analysis explained above for $\omega_T = 0 \ldots 0.1$ with steps of 0.0025 (Fig. 4.5B). To explore the relation between the coding level f and the range of decision-making, the optimal value of ω_T was determined for 11 different values of f from 0.1 to 0.225.

When changing the coding level, the network connectivities have to be adapted (Fig. 4.5C) to keep the up-state fixed-point firing rates at the

same values. Thus, the connectivity ω_+ was adjusted in steps of 0.0025 until the up-state fixed points matched the values for the parameters of the spiking simulation ($f = 0.2$ and $\omega_+ = 1.48$) (Fig. 4.5A).

For the final values of ω_+ for the respective coding levels, the up-state fixed-point values deviated by less than 2 Hz in the range of external inputs from 20 to 60 Hz. Note, that with changing ω_+ also ω_- changes, because of the normalization condition. In Fig. 4.5D the optimal range of decision-making was plotted against the coding level and fitted by a linear function.

Table A.3 Multiple-choice attractor model for primate data

A	Model Summary	
Populations	Six	
Topology	One module	
Connectivity	Full connectivity, no synaptic delay	
Neuron model	Leaky integrate-and-fire neurons, fixed voltage threshold, fixed absolute refractory periods	
Channel models	-	
Synapse models	Conductance-based synapses, AMPA and GABA$_A$ receptors (instantaneous rise, exponential decay), voltage-dependent NMDA receptors (exponential rise and decay)	
Plasticity	-	
Input	Independent fixed-rate poisson spike trains to all neurons	
Measurements	Spike activity	

B	Populations		
Total number of neurons	$N = 2,000$	Excitatory neurons	$N_E = 0.8 \cdot N$
		Inhibitory neurons	$N_I = 0.2 \cdot N$

Name	Size	Name	Size
Selective pool 1-4	$N_{S1-4} = f \cdot N_E$	Nonselective	$(1 - 4f) \cdot N_E$
		Inhibitory	$0.2 \cdot N$

C		Connectivity			
Source	**Target**	**Weight**	**Source**	**Target**	**Weight**
inhibitory $\mapsto$	all	$\omega_I = 1.125$	nonselective $\mapsto$	selective	$\omega_- = 0.8725$
excitatory $\mapsto$	inhibitory	$\omega = 1$	excitatory $\mapsto$	nonselective	$\omega = 1$
selective i $\mapsto$	selective i	$\omega_+ = 1.48$	selective 1 $\mapsto$	selective 3	$\omega_- = 0.8725$
selective 1+3 $\mapsto$	selective 2+4	$\omega_- + \omega_T = 0.8875$	selective 2 $\mapsto$	selective 4	$\omega_- = 0.8725$
selective 2+4 $\mapsto$	selective 1+3	$\omega_- + \omega_T = 0.8875$	selective 3 $\mapsto$	selective 1	$\omega_- = 0.8725$
			selective 4 $\mapsto$	selective 2	$\omega_- = 0.8725$

D	Neuron and Synapse Model
See Table A.1.	

E	Input	
Type	**Description**	
Poisson generator	Fixed rate $\nu = \nu_{ext} + \nu_n$, $\nu_{ext} = 2.4$ kHz, $\nu_n =$ sel. inputs with $n = 1\text{-}4$, one generator per neuron	

F	Measurements
Spike activity: firing-rates were calculated using the spike count in a 50 ms time window shifted by 5 ms steps and dividing by the number of neurons in the population and by the window size.	

Table A.4 Parameter set of multiple-choice attractor model for primate data

Parameter	Value	Parameter	Value
N_E	1,600	N_I	400
V_E	0 mV	V_I	-70 mV
V_L	-70 mV	N_{ext}	800
V_{thr}	-50 mV	ν_{ext}	2.4 kHz
V_{reset}	-55 mV	f	0.20
ω_+	1.48	ω_I	1.125
ω_-	0.8725	ω_T	0.015
C_m (excitatory)	0.5 nF	C_m (inhibitory)	0.2 nF
g_m (excitatory)	25 nS	g_m (inhibitory)	20 nS
τ_{ref} (excitatory)	2 ms	τ_{ref} (inhibitory)	1 ms
$g_{AMPA,ext}$ (excitatory)	2.08 nS	$g_{AMPA,ext}$ (inhibitory)	1.62 nS
$g_{AMPA,rec}$ (excitatory)	0.104 nS	$g_{AMPA,rec}$ (inhibitory)	0.081 nS
g_{GABA} (excitatory)	0.625 nS	g_{GABA} (inhibitory)	0.4865 nS
g_{NMDA} (excitatory)	0.14715 nS	g_{NMDA} (inhibitory)	0.1161 nS
$\tau_{NMDA,decay}$	100 ms	$\tau_{NMDA,rise}$	2 ms
τ_{AMPA}	2 ms	τ_{GABA}	10 ms
α	0.5 ms^{-1}		

A.1.6 Specifications of multiple-choice model for changes of mind

The multiple-choice attractor model used to fit and explain our psychophysical findings on changes of mind in a 2- and 4-choice RDM task (Chapter 5), was constructed by merging the relevant features of the binary changes of mind model (Chapter 3) with the multiple-choice model that accounted for the monkey data (Chapter 4). The general network structure is similar to the 2- and 4-choice model described in the last section, except that there is no spatial connectivity component ($\omega_T = 0$).

All default simulation parameters are listed in Table A.6. Compared to the original studies (Brunel and Wang, 2001; Wang, 2002), the AMPA/NMDA ratio was increased according to Eq. A.10 with $\delta = 0.08$.

1,000 trials of 3,500 ms with different random seeds were run for each parameter set and motion coherence.

Table A.5 Multiple-choice attractor model for changes of mind. A, and D-F as in Table A.3.

B				Populations		
Total number of neurons	$N = 500$			Excitatory neurons		$N_E = 0.8 \cdot N$
				Inhibitory neurons		$N_I = 0.2 \cdot N$

Name	Size	Name	Size
Selective pool 1 (up-right)	$N_{S1} = f \cdot N_E$	Nonselective	$(1 - 4f) \cdot N_E$
Selective pool 2 (up-left)	$N_{S2} = f \cdot N_E$	Inhibitory	$0.2 \cdot N$
Selective pool 3 (down-right)	$N_{S3} = f \cdot N_E$		
Selective pool 4 (down-left)	$N_{S4} = f \cdot N_E$		

C				Connectivity			
Source		**Target**	**Weight**	**Source**		**Target**	**Weight**
inhibitory	$\mapsto$	all	$\omega_I = 1.125$	nonselective	$\mapsto$	selective	$\omega_- = 0.88$
excitatory	$\mapsto$	inhibitory	$\omega = 1$	selective i	$\mapsto$	selective j	$\omega_- = 0.88$
excitatory	$\mapsto$	nonselective	$\omega = 1$	selective i	$\mapsto$	selective i	$\omega_+ = 1.48$

Table A.6 Parameter set of multiple-choice model for changes of mind

Parameter	Value	Parameter	Value
N_E	400	N_I	100
V_E	0 mV	V_I	-70 mV
V_L	-70 mV	N_{ext}	800
V_{thr}	-50 mV	ν_{ext}	2.4 kHz
V_{reset}	-55 mV	f	0.20
ω_+	1.48	ω_I	1.125
ω_-	0.88	α	0.5 ms^{-1}
C_m (excitatory)	0.5 nF	C_m (inhibitory)	0.2 nF
g_m (excitatory)	25 nS	g_m (inhibitory)	20 nS
τ_{ref} (excitatory)	2 ms	τ_{ref} (inhibitory)	1 ms
$g_{AMPA,ext}$ (excitatory)	2.08 nS	$g_{AMPA,ext}$ (inhibitory)	1.62 nS
$g_{AMPA,rec}$ (excitatory)	0.312 nS	$g_{AMPA,rec}$ (inhibitory)	0.243 nS
g_{GABA} (excitatory)	2.5 nS	g_{GABA} (inhibitory)	1.946 nS
g_{NMDA} (excitatory)	0.6213 nS	g_{NMDA} (inhibitory)	0.4902 nS
$\tau_{NMDA,decay}$	100 ms	$\tau_{NMDA,rise}$	2 ms
τ_{AMPA}	2 ms	τ_{GABA}	10 ms

A.2 Numerical simulation and data analysis

For each simulated trial in the full spiking network, the coupled differential equations that describe the dynamics of all cells and synapses (Eq. A.1-9) were integrated numerically using a second-order Runge-Kutta routine with a time-step of 0.02 ms.

Population firing rates were calculated by counting all spikes over a 50 ms window and dividing this sum by the number of neurons in the population and the window size. The time window was shifted with a time step of 5 ms.

The numerical integration of the mean-field equations was performed using either the Euler method, for the mean-field analysis presented in Chapter 4, or a second-order Runge-Kutta routine, for Chapters 3 and 5, with a time-step of 0.1 ms.

Both, the mean-field analysis described below and the spiking simulations were implemented in custom-made C++ programs (available upon request). Custom-made MATLAB programs were used for later analysis, fits of the simulation results and the numerical integration of the diffusion model (A.1.4).

Wherever present, error bars denote SEM over all correct trials for simulated reaction times. In the case of probabilities for correct choice and changes of mind the theoretically estimated SEM was calculated according to $\sqrt{p(1-p)/n}$ with n = 1,000 trials.

A.2.1 Fits to simulated behavioral data for binary changes of mind

Psychometric functions (Fig. 3.2, left panel) were fitted by a logistic function:

$$\text{Fraction correct} = (1 + \exp(-(\alpha + \beta \times coh)))^{-1}, \tag{A.34}$$

with motion coherence *coh*, and α and β as free parameters. The reaction time curve (Fig. 3.2, middle panel) was fitted by:

$$RT = \frac{A}{k \times coh} \tanh(Ak \times coh) + t_R, \tag{A.35}$$

with the free parameters A, k and t_R.

A.2.2 Fits to simulated "primate" multiple-choice data

The psychometric functions shown in Fig. 4.2 were fitted by a Weibull function:

$$\text{Fraction correct} = 1 - (1 - C) \times exp(-(coh/\alpha)^{\beta}),$$ (A.36)

where C is set as the chance level (0.25 and 0.5 for two and four alternatives, respectively), motion coherence coh, and α and β as free parameters. Reaction times were fitted as in Eq. A.35.

A.3 Detailed experimental paradigm of multiple-choice changes of mind

A.3.1 Experimental setup

a) Human subjects

Fifteen healthy young adults (10 female; mean age 22, range 19-27), right-handed and with normal vision participated in this study. None of the participants had any previous experience with visual psychophysics. Each participant underwent four experimental sessions during one day.

b) Experimental setup

The subjects sat in a dark room in front of a 21 inch flat-screen cathode ray tube video monitor (Sony Trinitron Multiscan CPD-G520 21). Viewing distance from the computer screen was 40 cm (Fig. 1A). Participants placed their head on a chin-and forehead-rest, which was calibrated for each participant before each experimental session. The visual stimuli were generated and data were collected using MATLAB (Mathworks, Natick, MA) and the Psychophysics toolbox-3 (Brainard, 1997; Kleiner et al., 2007) on an ASUS P5K SE/EPU computer running Microsoft Windows XP at a frame rate of 75 Hz.

c) Visual stimuli

The visual stimuli presented in this task were constructed in the following way. The RDM stimulus consisted of a multi-component pattern of small white moving dots (small filled squares with an edge length of 2 pixels). The dots appeared within a circular aperture (diameter = 5.0º) at the center of the screen. The percept of apparent dot-motion was created as follows: The stimulus consisted of three independent streams of dots that were alternately presented every three frames. With the next presentation of a particular set of dots three frames later, these dots were displaced as follows. Dependent on the respective amount of coherent motion, a certain percentage of dots were shifted in a particular direction, while the other dots were replaced at a random location. In that way, the dot positions of frame 3 for example were correlated with those of frame 6, but not with frame one, two, four, or five (Palmer et al., 2005; Roitman and Shadlen, 2002; Shadlen and Newsome, 2001). The coherently moving

dots had a speed of 6.0º/s. Dot density was 16.7 dots/(deg^2·s). We used a set of eight different coherence levels (0%, 3.2%, 6.4%, 12.8%, 25.6%, 51.2%, 76.8% and 100%).

The R-targets were constructed as yellow circles (diameter = 2.0º) located at the corners of a virtual square around the central fixation-mark (edge length 28º, and thus 19.8º distance to the center). The location of the R-targets indicated the possible directions of coherent motion in each trial. They could appear either:

a) in each of the four corners of the virtual square (4-choice trials), or

b) in just two of the four corners (2-choice trials).

The two presented R-targets could be either 180º apart and thus symmetrically located with respect to the fixation-mark, or 90º apart, and therefore contiguous on one of the virtual square's sides[20] (Fig. 1B, top and bottom left panel respectively). The R-targets remained present on the screen until the end of the trial.

d) Feedback

After each trial, participants received visual feedback (*"error"*, *"good"*) according to their performance. They also received a *"time out"* alert message whenever they exceeded time-out #1 or #2. *"wait for cue!"* appeared if the mouse pointer was moved out of the start-target before the RDM stimulus had appeared. Finally, *"no target hit!"* was displayed every time participants were not accurate in selecting the start target or an R-target area properly. These trials were excluded.

A.3.2 Detection of changes of mind

The criteria for a change of mind were: (1) the area between the mouse trajectory and either the horizontal, or vertical axis had to exceed 2100 pixel^2 in a screen quadrant other than the finally selected. (2) The trajectory excursion had to exceed the start-target area by 20 pixels in a quadrant other than the finally chosen. This quadrant then denoted the initial choice. We further rechecked the accuracy of the selection algorithm by visual inspection.

---

[20] In the 90º-case, the second adjacent target was chosen randomly as +/- 90º from the correct R-target.

Bibliography

Abeles M (1991) Corticonics: Neural Circuits of the Cerebral Cortex. New York: Cambridge University Press.

Albantakis L, Deco G (2009) The encoding of alternatives in multiple-choice decision making. Proc Natl Acad Sci U S A 106:10308-10313.

Albantakis L, Deco G (2011) Changes of mind in an attractor network of decision-making. PLoS Comput Biol 7:e1002086.

Amit DJ (1995) The Hebbian paradigm reintegrated: local reverberations as internal representations. Behav Brain Sci 18:617-626.

Amit DJ, Brunel N (1997) Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. Cereb Cortex 7:237-252.

Andersen RA (1995) Encoding of intention and spatial location in the posterior parietal cortex. Cereb Cortex 5:457-469.

Andersen RA, Cui H (2009) Intention, Action Planning, and Decision Making in Parietal-Frontal Circuits. Neuron 63:568-583.

Beck JM, Ma WJ, Kiani R, Hanks T, Churchland AK, Roitman J, Shadlen MN, Latham PE, Pouget A (2008) Probabilistic population codes for Bayesian decision making. Neuron 60:1142-1152.

Ben-Yishai R, Bar-Or RL, Sompolinsky H (1995) Theory of orientation tuning in visual cortex. Proc Natl Acad Sci U S A 92:3844-3848.

Bogacz R, Gurney K (2007) The basal ganglia and cortex implement optimal decision making between alternative actions. Neural Comput 19:442-477.

Bogacz R, Usher M, Zhang J, McClelland JL (2007) Extending a biologically inspired model of choice: multi-alternatives, nonlinearity and value-based multidimensional choice. Philos T R Soc B 362:1655-1670.

Bogacz R, Wagenmakers EJ, Forstmann BU, Nieuwenhuis S (2010) The neural basis of the speed–accuracy tradeoff. Trends Neurosci 33:10-16.

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

Brainard DH (1997) The Psychophysics Toolbox. Spat Vis 10:433-436.

Britten KH, Shadlen MN, Newsome WT, Movshon JA (1992) The analysis of visual motion: a comparison of neuronal and psychophysical performance. J Neurosci 12:4745-4765.

Britten KH, Shadlen MN, Newsome WT, Movshon JA (1993) Responses of neurons in macaque MT to stochastic motion signals. Vis Neurosci 10:1157-1169.

Brown S, Heathcote A (2008) The simplest complete model of choice response time: Linear ballistic accumulation. Cognitive Psychol 57:153-178.

Brunel N (2003) Dynamics and plasticity of stimulus-selective persistent activity in cortical network models. Cereb Cortex 13:1151-1161.

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

Busemeyer JR, Townsend JT (1993) Decision field theory: a dynamic-cognitive approach to decision making in an uncertain environment. Psychol Rev 100:432-459.

Churchland AK, Kiani R, Shadlen MN (2008) Decision-making with multiple alternatives. Nat Neurosci 11:693-702.

Churchland AK, Kiani R, Chaudhuri R, Wang XJ, Pouget A, Shadlen MN (2011) Variance as a signature of neural computations during decision making. Neuron 69:818-831.

Cisek P, Puskas GA, El-Murr S (2009) Decisions in changing conditions: the urgency-gating model. J Neurosci 29:11560-11571.

Cui H, Andersen RA (2007) Posterior Parietal Cortex Encodes Autonomously Selected Motor Plans. Neuron 56:552-559.

Deco G, Rolls ET (2005) Attention, short-term memory, and action selection: a unifying theory. Prog Neurobiol 76:236-256.

Deco G, Rolls ET (2006) Decision-making and Weber's law: a neurophysiological model. Eur J Neurosci 24:901-916.

Deco G, Romo R (2008) The role of fluctuations in perception. Trends Neurosci 31:591-598.

Deco G, Rolls ET, Romo R (2009) Stochastic dynamics as a principle of brain function. Prog Neurobiol 88:1-16.

Deco G, Rolls ET, Romo R (2010) Synaptic dynamics and decision making. Proc Natl Acad Sci U S A 107:7545-7549.

Deco G, Jirsa VK, McIntosh AR (2011) Emerging concepts for the dynamical organization of resting-state activity in the brain. Nat Rev Neurosci 12:43-56.

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

Ditterich J (2006b) Stochastic models of decisions about motion direction: Behavior and physiology. Neural Networks 19:981-1012.

Ditterich J (2010) A Comparison between Mechanisms of Multi-Alternative Perceptual Decision Making: Ability to Explain Human Behavior, Predictions for Neurophysiology, and Relationship with Decision Theory. Front Neurosci 4:184.

Ditterich J, Mazurek ME, Shadlen MN (2003) Microstimulation of visual cortex affects the speed of perceptual decisions. Nat Neurosci 6:891-898.

Djurfeldt M, Ekeberg O, Lansner A (2008a) Large-scale modeling - a tool for conquering the complexity of the brain. Front Neuroinform 2:1.

Djurfeldt M, Lundqvist M, Johansson C, Rehn M, Ekeberg O, Lansner A (2008b) Brain-scale simulation of the neocortex on the IBM Blue Gene/L supercomputer. IBM J Res Dev 52:31-41.

Dragalin VP (1999) Multihypothesis sequential probability ratio tests-Part I : Asymptotic optimality. IEEE T Inform Theory 45:2448-2461.

Furman M, Wang XJ (2008) Similarity Effect and Optimal Control of Multiple-Choice Decision Making. Neuron 60:1153-1168.

Gardiner CW (1985) Handbook of stochastic methods (2nd ed.) New York: Springer.

Gerstner W, Kistler WM (2002) Spiking Neuron Models. Single Neurons, Populations, Plasticity. Ney York: Cambridge University Press.

Gerstner W, Naud R (2009) Neuroscience. How good are neuron models? Science 326:379-380.

Gnadt JW, Andersen RA (1988) Memory related motor planning activity in posterior parietal cortex of macaque. Exp Brain Res 70:216-220.

Gold JI, Shadlen MN (2001) Neural computations that underlie decisions about sensory stimuli. Trends Cogn Sci 5:10-16.

Gold JI, Shadlen MN (2002) Banburismus and the brain: decoding the relationship between sensory stimuli, decisions, and reward. Neuron 36:299-308.

Gold JI, Shadlen MN (2007) The Neural Basis of Decision Making. Annu Rev Neurosci 30:535-574.

Grabenhorst F, Rolls ET (2011) Value, pleasure and choice in the ventral prefrontal cortex. Trends Cogn Sci 15:56-67.

Green DM, Swets JA (1966) Signal detection theory and psychophysics. New York: John Wiley & Sons, Inc

Heekeren HR, Marrett S, Ungerleider LG (2008) The neural systems that mediate human perceptual decision making. Nat Rev Neurosci 9:467-479.

Heekeren HR, Marrett S, Ruff DA, Bandettini PA, Ungerleider LG (2006) Involvement of human left dorsolateral prefrontal cortex in perceptual decision making is independent of response modality. Proc Natl Acad Sci U S A 103:10023-10028.

Hernandez A, Nacher V, Luna R, Zainos A, Lemus L, Alvarez M, Vazquez Y, Camarillo L, Romo R (2010) Decoding a perceptual decision process across cortex. Neuron 66:300-314.

Hestrin S, Sah P, Nicoll RA (1990) Mechanisms generating the time course of dual component excitatory synaptic currents recorded in hippocampal slices. Neuron 5:247-253.

Hick WE (1952) On the rate of gain of information. Q J Exp Psycho 4:11-26.

Hill CM (1898) On Choice. Am J Psychol 9:587-590.

Horwitz GD, Newsome WT (1999) Separate signals for target selection and movement specification in the superior colliculus. Science 284:1158-1161.

Huk AC, Shadlen MN (2005) Neural activity in macaque parietal cortex reflects temporal integration of visual motion signals during perceptual decision making. J Neurosci 25:10420-10436.

Izhikevich EM, Edelman GM (2008) Large-scale model of mammalian thalamocortical systems. Proc Natl Acad Sci U S A 105:3593-3598.

Jahr CE, Stevens CF (1990) Voltage dependence of NMDA-activated macroscopic conductances predicted by single-channel kinetics. J Neurosci 10:3178-3182.

Kiani R, Shadlen MN (2009) Representation of Confidence Associated with a Decision by Neurons in the Parietal Cortex. Science 324:759-764.

Kiani R, Hanks TD, Shadlen MN (2008) Bounded Integration in Parietal Cortex Underlies Decisions Even When Viewing Duration Is Dictated by the Environment. J Neurosci 28:3017-3029.

Kim JN, Shadlen MN (1999) Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. Nat Neurosci 2:176-185.

Kleiner M, Brainard D, Pelli D (2007) What's new in psychtoolbox-3? Perception 36:ECVP Abstract Supplement.

Laming DRJ (1968) Information theory of choice-reaction times. London: Academic P.

Lansner A (2009) Associative memory models: from the cell-assembly theory to biophysically detailed cortex simulations. Trends Neurosci 32:178-186.

Leite FP, Ratcliff R (2009) Modeling reaction time and accuracy of multiple-alternative decisions. Atten Percept Psycho 72:246-273.

Lo CC, Wang XJ (2006) Cortico–basal ganglia circuit mechanism for a decision threshold in reaction time tasks. Nat Neurosci 9:956-963.

Loh M, Rolls ET, Deco G (2007) A dynamical systems hypothesis of schizophrenia. PLoS Comput Biol 3:e228.

Luce R (1991) Response Times: Their Role in Inferring Elementary Mental Organization: Oxford University Press, USA.

Lund JS, Angelucci A, Bressloff PC (2003) Anatomical substrates for functional columns in macaque monkey primary visual cortex. Cereb Cortex 13:15-24.

Ma WJ, Beck JM, Latham PE, Pouget A (2006) Bayesian inference with probabilistic population codes. Nat Neurosci 9:1432-1438.

Machens CK, Romo R, Brody CD (2005) Flexible control of mutual inhibition: a neural model of two-interval discrimination. Science 307:1121-1124.

Markram H (2006) The blue brain project. Nat Rev Neurosci 7:153-160.

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

Masquelier T, Albantakis L, Deco G (2011) The timing of vision - how neural processing links to different temporal dynamics. Front Psychol 2:151.

Mazurek ME, Roitman JD, Ditterich J, Shadlen MN (2003) A role for neural integrators in perceptual decision making. Cereb Cortex 13:1257-1269.

McMillen T, Holmes P (2006) The dynamics of choice among multiple alternatives. J Math Psychol 50:30-57.

Mongillo G, Barak O, Tsodyks M (2008) Synaptic Theory of Working Memory. Science 319:1543-1546.

Neyman J, Pearson ES (1933) On the Problem of the Most Efficient Tests of Statistical Hypotheses. Philos T R Soc B 231:289-337.

Niwa M, Ditterich J (2008) Perceptual Decisions between Multiple Directions of Visual Motion. J Neurosci 28:4435-4445.

Opris I, Bruce C (2005) Neural circuitry of judgment and decision mechanisms. Brain Res Rev 48:509-526.

Palmer J, Huk AC, Shadlen MN (2005) The effect of stimulus strength on the speed and accuracy of a perceptual decision. J Vis 5:1-1.

Platt ML, Glimcher PW (1999) Neural correlates of decision variables in parietal cortex. Nature 400:233-238.

Pleskac TJ, Busemeyer JR (2010) Two-stage dynamic signal detection: a theory of choice, decision time, and confidence. Psychol Rev 117:864-901.

Purcell BA, Heitz RP, Cohen JY, Schall JD, Logan GD, Palmeri TJ (2010) Neurally constrained modeling of perceptual decision making. Psychol Rev 117:1113-1143.

Ratcliff R (1978) A theory of memory retrieval. Psychol Rev 85:59-108.

Ratcliff R, Rouder JN (1998) Modeling response times for two-choice decisions. Psychol Sci 9:347-356.

Ratcliff R, Smith PL (2004) A Comparison of Sequential Sampling Models for Two-Choice Reaction Time. Psychol Rev 111:333-367.

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

Ratcliff R, Cherian A, Segraves M (2003) A comparison of macaque behavior and superior colliculus neuronal activity to predictions from models of two-choice decisions. J Neurophysiol 90:1392-1407.

Reddi BA, Carpenter RH (2000) The influence of urgency on decision time. Nat Neurosci 3:827-830.

Renart A, Brunel N, Wang XJ (2003) Mean-field theory of recurrent cortical networks: from irregularly spiking neurons to working memory. In: Computational neuroscience: a comprehensive approach. (Feng J, ed), pp 431-490. Boca Raton, FL: Chapman & Hall/CRC.

Resulaj A, Kiani R, Wolpert DM, Shadlen MN (2009) Changes of mind in decision-making. Nature 461:263-266.

Roitman JD, Shadlen MN (2002) Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. J Neurosci 22:9475-9489.

Rolls ET, Webb TJ (2011) Cortical attractor network dynamics with diluted connectivity. Brain Res In Press, Accepted Manuscript.

Rolls ET, Loh M, Deco G (2008a) An attractor hypothesis of obsessive-compulsive disorder. Eur J Neurosci 28:782-793.

Rolls ET, Loh M, Deco G, Winterer G (2008b) Computational models of schizophrenia and dopamine modulation in the prefrontal cortex. Nat Rev Neurosci 9:696-709.

Romo R, Salinas E (2003) Flutter discrimination: neural codes, perception, memory and decision making. Nat Rev Neurosci 4:203-218.

Roxin A, Ledberg A (2008) Neurobiological models of two-choice decision making can be reduced to a one-dimensional nonlinear diffusion equation. PLoS Comput Biol 4:e1000046.

Sakai Y, Funahashi S, Shinomoto S (1999) Temporally correlated inputs to leaky integrate-and-fire models can reproduce spiking statistics of cortical neurons. Neural Networks 12:1181-1190.

Salin PA, Prince DA (1996) Spontaneous GABAA receptor-mediated inhibitory currents in adult rat somatosensory cortex. J Neurophysiol 75:1573-1588.

Salzman CD, Murasugi CM, Britten KH, Newsome WT (1992) Microstimulation in visual area MT: effects on direction discrimination performance. J Neurosci 12:2331-2355.

Schall JD (2001) Neural basis of deciding, choosing and acting. Nat Rev Neurosci 2:33-42.

Schall JD (2003) Neural correlates of decision processes: neural and mental chronometry. Curr Opin Neurobiol 13:182-186.

Shadlen MN, Newsome WT (1996) Motion perception: seeing and deciding. Proc Natl Acad Sci U S A 93:628-633.

Shadlen MN, Newsome WT (2001) Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. J Neurophysiol 86:1916-1936.

Smith P, Ratcliff R (2004) Psychology and neurobiology of simple decisions. Trends Neurosci 27:161-168.

Snyder LH, Batista AP, Andersen RA (1997) Coding of intention in the posterior parietal cortex. Nature 386:167-170.

Soltani A, Wang X-J (2009) Synaptic computation underlying probabilistic inference. Nat Neurosci 13:112-119.

Sporns O, Tononi G, Kotter R (2005) The human connectome: A structural description of the human brain. PLoS Comput Biol 1:e42.

Spruston N, Jonas P, Sakmann B (1995) Dendritic glutamate receptor channels in rat hippocampal CA3 and CA1 pyramidal neurons. J Physiol 482 (Pt 2):325-352.

Stone M (1960) Models for choice-reaction time. Psychometrika 25:251-260.

Tanner WP, Jr., Swets JA (1954) A decision-making theory of visual detection. Psychol Rev 61:401-409.

Townsend JT, Ashby FG (1983) The stochastic modeling of elementary psychological processes. Cambridge, UK: Cambridge University Press.

Tuckwell H (1988) Introduction to Theoretical Neurobiology. Cambridge, UK: Cambridge University Press.

Usher M, McClelland JL (2001) The time course of perceptual choice: the leaky, competing accumulator model. Psychol Rev 108:550-592.

Usher M, Olami Z, McClelland JL (2002) Hick's Law in a Stochastic Race Model with Speed–Accuracy Tradeoff. J Math Psychol 46:704-715.

Vandekerckhove J, Tuerlinckx F (2007) Fitting the Ratcliff diffusion model to experimental data. Psychon Bull Rev 14:1011-1026.

Vickers D (1970) Evidence for an accumulator model of psychophysical discrimination. Ergonomics 13:37-58.

Vickers D (1979) Decision processes in visual perception. New York ; London: Academic Press.

Wald A (1947) Sequential Analysis. New York: Wiley.

Wang XJ (2001) Synaptic reverberation underlying mnemonic persistent activity. Trends Neurosci 24:455-463.

Wang XJ (2002) Probabilistic decision making by slow reverberation in cortical circuits. Neuron 36:955-968.

Wang XJ (2008) Decision Making in Recurrent Neuronal Circuits. Neuron 60:215-234.

Webb TJ, Rolls ET, Deco G, Feng J (2011) Noise in Attractor Networks in the Brain Produced by Graded Firing Rate Representations. PLoS ONE 6:e23630.

Wong KF, Wang XJ (2006) A Recurrent Network Mechanism of Time Integration in Perceptual Decisions. J Neurosci 26:1314-1328.

Wong KF, Huk AC (2008) Temporal dynamics underlying perceptual decision making: Insights from the interplay between an attractor model and parietal neurophysiology. Front Neurosci 2:245-254.

Wong KF, Huk AC, Shadlen MN, Wang XJ (2007) Neural circuit dynamics underlying accumulation of time-varying evidence during perceptual decision making. Front Comput Neurosci 1:6.

Xiang Z, Huguenard JR, Prince DA (1998) GABAA receptor-mediated currents in interneurons and pyramidal cells of rat visual cortex. J Physiol 506 (Pt 3):715-730.

List of Abbreviations

2AFC	2-alternative forced-choice
AMPA	α-amino-3-hydroxy-5-methyl-4-isoazoleproprionic acid
DDM	drift diffusion model
dlPFC	dorso-lateral prefrontal cortex
EPSC	excitatory post-synaptic current
ER	error rate
FEF	frontal eye-field
FF	feedforward
FFI	feedforward inhibition
GABA	γ-aminobutyric acid
IPSC	inhibitory post-synaptic current
LCA	leaky competing accumulator
LIF	leaky integrate-and-fire
LIP	lateral interaparietal cortex
LR	likelihood ratio
MT	middle temporal area
NMDA	N-methyl-D-aspartate acid
O-U	Ornstein-Uhlenbeck
PDF	probability density function
PFC	prefrontal cortex
PPC	posterior parietal cortex
PRR	parietal reach region
RDM	random-dot motion
RF	response field
R-target	response target
RT	reaction time
SAT	speed-accuracy tradeoff
SC	superior colliculus
SDT	signal detection theory
SPRT	sequential probability ratio test