


**Provided for non-commercial research and educational use only.
Not for reproduction, distribution or commercial use.**

This chapter was originally published in the book *Handbook of Reward and Decision Making*, published by Elsevier, and the attached copy is provided by Elsevier for the author's benefit and for the benefit of the author's institution, for non-commercial research and educational use including without limitation use in instruction at your institution, sending it to specific colleagues who know you, and providing a copy to your institution's administrator.



All other uses, reproduction and distribution, including without limitation commercial reprints, selling or licensing copies or access, or posting on open internet sites, your personal or institution's website or repository, are prohibited. For exceptions, permission may be sought for such use through Elsevier's permissions site at:

<http://www.elsevier.com/locate/permissionusematerial>

From Rafal Bogacz, Optimal decision-making theories. In: Dr. Jean-Claude Dreher and Léon Tremblay, editors, *Handbook of Reward and Decision Making*. Oxford: Academic Press, 2009, pp. 375-397.

ISBN: 978-0-12-374620-7

© Copyright 2009 Elsevier Inc.
Academic Press

Author's personal copy

Part Five

Computational Models of the Reward System and Decision Making

18 Optimal decision-making theories

Rafal Bogacz

Department of Computer Science, University of Bristol, Bristol BS8 1UB, United Kingdom

Abstract

In case of many decisions based on sensory information, the sensory stimulus or its neural representation are noisy. This chapter reviews theories proposing that the brain implements statistically optimal strategies for decision making on the basis of noisy information. These strategies maximize the accuracy and speed of decisions, as well as the rate of receiving rewards for correct choices. The chapter first reviews computational models of cortical decision circuits that can optimally perform choices between two alternatives. Then, it describes a model of cortico-basal-ganglia circuit that implements the optimal strategy for choice between multiple alternatives. Finally, it shows how the basal ganglia may modulate decision processes in the cortex, allowing cortical neurons to represent the probabilities of alternative choices being correct. For each set of theories their predictions are compared with existing experimental data.

Key points

1. Integrating sensory information over time increases the accuracy of decisions.
2. Optimal decision strategies describe when to stop the integration of information, and they minimize the decision time for any required accuracy.
3. Several models of cortical decision circuits have been proposed that implement the optimal strategy for choice between two alternatives.
4. It has been proposed that the optimal strategy for choice between multiple alternatives is implemented in the cortico-basal-ganglia circuit.
5. In the proposed model, the cortico-basal-ganglia circuit computes the posterior probabilities of alternatives being correct, given the sensory evidence.

18.1 Introduction

Imagine an animal trying to decide if a shape moving behind the leaves is predator or prey. The sensory information the animal receives may be noisy and partial (e.g., occluded by the leaves) and hence need to be accumulated over time to gain sufficient accuracy, but the speed of such choices is also of critical importance. Due to the noisy nature of the sensory information, such decisions can be considered as statistical problems. This chapter reviews theories assuming that during perceptual decisions the brain performs statistically optimal tests, thereby maximizing the speed and accuracy of choices.

The optimal decision-making theories are motivated by an assumption that evolutionary pressure has been promoting animals that make fast and accurate choices. Although optimality is not always achieved, these theories provide interpretation for existing data and further experimental predictions that can guide empirical research. The optimal decision-making strategies are often precisely defined and relatively simple, thus they constrain the parameters of the models of neural decision circuits.

This chapter is organized as follows. [Section 18.2](#) reviews models assuming that cortical decision circuits implement an optimal test for choice between two alternatives. [Section 18.3](#) reviews the model of cortico-basal-ganglia circuit assuming that it implements optimal choice between multiple alternatives; this model is closely related to that described in [Chapter 19](#). [Section 18.4](#) shows how the basal ganglia may modulate the integration of sensory evidence in the cortex and allow cortical neurons represent the probabilities of alternative choices being correct. [Section 18.5](#) discusses open questions. In addition, the Appendices contain essential derivations presented step-by-step (without skipping any calculations) such that they can be understood without extensive mathematical background.

In this chapter we focus on models of decision making in highly practiced task; the models describing task acquisition are discussed in [Chapter 19](#). Throughout this chapter we use population level models describing activities of neuronal populations rather than individual neurons. Although this level of modeling is unable to capture many important details of neural decision circuits, it helps in understanding the essence of computation performed by various populations during decision process.

18.2 Models of optimal decision making in the cortex

In this section we first briefly review the responses of cortical neurons during a choice task; a more detailed review is available in [Chapter 8](#). Next, the statistically optimal decision strategy and its proposed neural implementations are described. Finally, the predictions of the models are compared with experimental data.

18.2.1 *Neurobiology of decision processes in the cortex*

The neural bases of decision are typically studied in the motion discrimination task, in which a monkey is presented with a display containing moving dots [1]. A fraction of the dots moves left on some trials or right on other trials, while the rest is moving randomly. The task of the animal is to make a saccade (i.e., an eye movement) in the direction of motion of the majority of dots.

One of the areas critically important for this task is the medial temporal (MT) area in the visual cortex. The neurons in area MT respond selectively for a particular directions of motion, thus the MT neurons preferring left motion are more active on trials when the majority of dots move left, while the neurons preferring right motion are more active on trials when the majority of dots move right [1]. However, their activity is also very noisy, as the stimulus itself is noisy.

Let us now consider the decision faced by the areas receiving input from MT. Let us imagine a judge listening to the activity from the two populations of MT neurons, each preferring one of the two directions. The MT neurons produce spikes that can be interpreted as votes for the alternative directions. The task of the judge is to choose the alternative receiving more votes, that is, corresponding to the MT population with higher mean firing rate. Note however that the judge cannot measure the mean firing

rate instantaneously, as the spikes are discrete events spread out over time. Thus the judge needs to observe the MT activity for a period of time and integrate the evidence or “count the votes” until it reaches a certain level of confidence.

Indeed, such an information integration process has been observed in the lateral intraparietal (LIP) area and frontal eye field during the motion coherence task. Neurons in area LIP respond selectively before and during saccades in particular directions. It has been observed that the LIP neurons selective for the chosen direction gradually increase their firing rate during the motion discrimination task [2, 3], and the data indicate that these neurons integrate the input from corresponding MT neurons over time [4–6]. Thus, using the judge metaphor, the LIP neurons “count the votes” and represent the total numbers of votes in their firing rate. The neurons representing integrated evidence have also been observed in other tasks in areas within the frontal lobe [7–10].

18.2.2 Optimal stopping criterion for two alternatives

The analysis still leaves an open question: When should the integration process be stopped and action executed? This subsection first considers this question intuitively, then it presents an optimal stopping criterion formally, and describes in what sense it is optimal.

Stopping criteria. The simplest possible criterion is to stop the integration whenever the integrated evidence, or the total number of votes, for one of the alternatives reaches a threshold. This strategy is known as the race model [11]. Another possibility is to stop the integration when the *difference* between the integrated evidence in favor of the winning and losing alternatives exceeds a threshold. This strategy is referred to as the diffusion model [12–14].

The diffusion model is usually formulated in a different but equivalent way: It includes a single integrator which accumulates the difference between the sensory evidence supporting the two alternatives. A choice is made in favor of the first alternative if the integrator exceeds a positive threshold, or in favor of the second alternative if the integrator decreases below a negative threshold.

As will be demonstrated formally here, the diffusion model provides an optimal stopping criterion. The advantage of diffusion over race models can also be seen intuitively, as the diffusion model allows the decision process to be modulated by the amount of conflict between evidence supporting alternatives on a given trial: Note that decisions will take longer when the evidence for the two alternatives is similar, and importantly, that decisions will be faster when there is little evidence for the losing alternative. By contrast, in the race model the decision time does not depend on the level of evidence for the losing alternative [15].

Statistical formulation. Let us now formalize the decision problem. For simplicity, let us assume that time can be divided into discrete steps. Let $x_i(t)$ denote the sensory evidence supporting alternative i at time t , which in the motion discrimination task corresponds to the firing rate of the MT neurons selective for alternative i at time t . Let y_i denote the evidence for alternative i integrated until time t :

$$y_i = \sum_{\tau=1}^t x_i(\tau). \quad (18.1)$$

Gold and Shadlen [16, 17] formulated the decision as a statistical problem in the following way: They assumed that $x_i(t)$ come from a normal distribution with mean μ_i and

standard deviation σ . They defined hypotheses H_i stating that the sensory evidence supporting alternative i has higher mean:

$$H_1: \mu_1 = \mu^+, \mu_2 = \mu^-; H_2: \mu_1 = \mu^-, \mu_2 = \mu^+, \tag{18.2}$$

where $\mu^+ > \mu^-$. The optimal procedure for distinguishing between these hypotheses is provided by the sequential probability ratio test (SPRT) [18], which is equivalent to the diffusion model [12], as will be shown. According to SPRT, at each moment of time t , one computes the ratio of the likelihoods of the sensory evidence given the hypotheses:

$$R = \frac{P(x(1..t)|H_1)}{P(x(1..t)|H_2)}, \tag{18.3}$$

where $x(1..t)$ denotes a set of all sensory evidence observed so far [i.e., $x_1(1), x_2(1), \dots, x_1(t), x_2(t)$]. If R exceeds a threshold Z_1 or decreases below a lower threshold Z_2 , the decision process is stopped and the choice is made (H_1 if $R > Z_1$ or H_2 if $R < Z_2$). Otherwise the decision process continues and another sample of sensory information is observed. Thus note that SPRT observes sensory evidence only for as long as it is necessary to distinguish between the hypotheses.

The relationship between SPRT and the diffusion model is described in [Appendix 18A](#). Its first part shows that R in the SPRT changes in a similar way as the single integrator in the diffusion model (which accumulates the difference between the sensory evidence supporting the two alternatives). Namely, if at a given moment of time t , the sensory evidence supports the first alternative more than the second [i.e., $x_1(t) > x_2(t)$] then R increases; and otherwise R decreases. The second part of [Appendix 18A](#) shows that $\log R$ is exactly proportional to the difference between the integrated evidence:

$$\log R = g(y_1 - y_2), \tag{18.4}$$

where g is a constant. Thus according to SPRT, the decision process will be stopped when the difference $(y_1 - y_2)$ exceeds a positive threshold (equal to $\log Z_1/g$) or decreases below a negative threshold (equal to $\log Z_2/g$). Hence the SPRT (with hypotheses of [Eq. 18.2](#)) is equivalent to the diffusion model.

Optimality. The SPRT is optimal in the following sense: it minimizes the average decision time for any required accuracy [19]. Let us illustrate this property by comparing the race and the diffusion models. In both models, speed and accuracy are controlled by the height of decision threshold, and there exists a speed-accuracy tradeoff. But if we choose the thresholds in both models giving the same accuracy for a given sensory input, then the diffusion model will on average be faster than the race model.

One can ask if this optimality property could bring benefits to animals. If we consider a scenario in which an animal receives rewards for correct choices, then making fast choices allows an animal to receive more rewards per unit of time. In particular, Gold and Shadlen [17] considered a task in which the animal receives a reward for correct choices, there is no penalty for errors, and there is a fixed delay D between the response and the onset of the next choice trial. In this task the reward rate is equal to the ratio of accuracy AC and the average duration of the trial [17]:

$$RR = \frac{AC}{RT + D}, \tag{18.5}$$

where RT denotes the average reaction time. The reward rate depends on decision threshold. But [Appendix 18B](#) shows that the diffusion model with appropriately chosen threshold gives higher or equal reward rate than any other decision strategy. Furthermore, the diffusion model maximizes reward rate not only for the task described above, but also for a wide range of other tasks [\[20\]](#).

18.2.3 Neural implementations

As mentioned in the previous subsection, the optimal decision strategy is to make a choice when the difference between the evidence supporting the two alternatives exceeds a threshold. However, the neurophysiological data suggest that an animal makes a choice when the activity of integrator neurons selective for a given alternative exceeds a fixed threshold [\[3, 21\]](#). One way to reconcile these two observations is to assume that the firing rate of these neurons represents the difference in evidence supporting the two alternatives. Several models have been proposed that exhibit this property, and they are briefly reviewed in this subsection.

Let us refer to the population of cortical integrator neurons selective for a particular alternative as an integrator. In the model shown in [Fig. 18.1A](#), the integrators directly accumulate the difference between the activities of sensory neurons via feed-forward inhibitory connections [\[2, 22\]](#).

Thus this model is equivalent to the diffusion model. An alternative model shown in [Fig. 18.1B](#) assumes that integrators mutually inhibit each other [\[23\]](#). An analysis of the dynamics of this model reveals that, for certain parameter values, the activity of integrators is approximately proportional to the difference between the integrated evidence, thus it also approximates the diffusion model [\[20\]](#).

In the models shown in [Fig. 18.1A and 18.1B](#), the sensory neurons send both excitatory and inhibitory connections. But in the brain the cortico-cortical connections are only excitatory, thus [Fig. 18.1C and 18.1D](#) shows more realistic versions of the models in [Fig. 18.1A and 18.1B](#) in which the inhibition is provided by a population of inhibitory inter-neurons (the model in [Fig. 18.1D](#) has been proposed in [\[24\]](#)). It is easy to set the parameters of the model shown in [Fig. 18.1C](#) so it also integrates the difference between sensory evidence (in particular, the weights of connections between sensory and integrator populations need to be twice as large as the other weights). Similarly Wong and Wang [\[25\]](#) have shown that the model in [Fig. 18.1D](#) can also for certain parameters integrate the difference between sensory evidence.

In summary, there are several architectures of cortical decision networks, shown in [Fig. 18.1](#), that can approximate the optimal decision making for two alternatives. Matlab codes allowing simulation and comparison of performance of the models described in this chapter are available <http://www.cs.bris.ac.uk/home/rafal/optimal/codes.html>

18.2.4 Comparison with experimental data

The diffusion model has been shown to fit the reaction time (RT) distributions from a wide range of choice tasks [\[26–28\]](#). Furthermore, the diffusion model has been shown to describe the patterns of RTs and the growth of information in neurons involved in the decision process better than the race model [\[29–31\]](#).

The diffusion model has been also used to explain the effect of stimulation of MT and LIP neurons during the motion discrimination task. Ditterich et al. [\[4\]](#) showed that stimulation of MT neurons selective for a particular alternative produced two effects: reduced

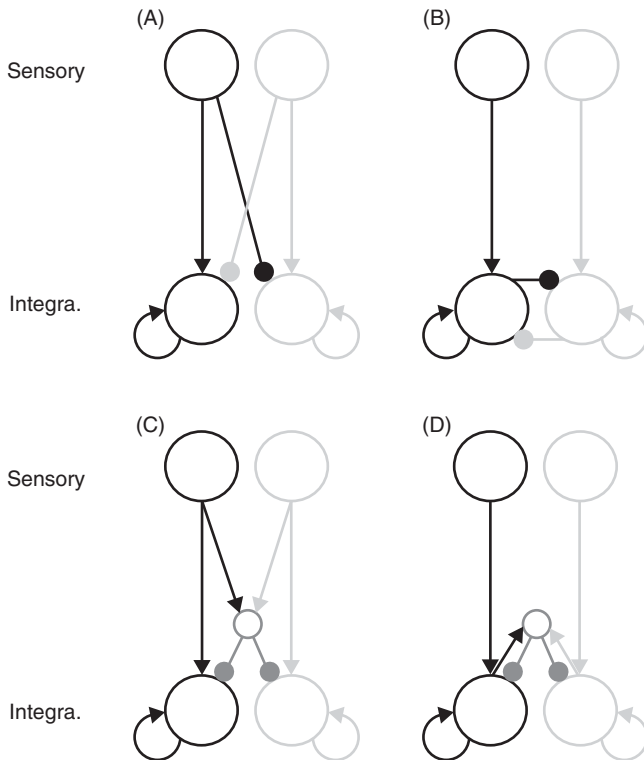


Figure 18.1 Connections between neuronal populations in cortical models of decision making (A) Shadlen and Newsome model, (B) Usher and McClelland model, (C) feed-forward pooled inhibition model, (D) Wang model. Open circles denote neuronal populations. Arrows denote excitatory connections; lines ended with circles denote inhibitory connections. Thus the open circles with self-excitatory connections denote the integrator populations. Black and gray pathways correspond to neuronal populations selective for the two alternative choices. Sensory – sensory cortex encoding relevant aspects of stimuli; Integra. – cortical region integrating sensory evidence.

RTs on trials when this alternative was chosen, and increased RTs when the other alternative was chosen. The second effect suggests that the decision is made on the basis of the difference between integrated evidence, and thus supports the diffusion model over the race model. Hanks et al. [32] compared these effects of MT stimulation with the effects of LIP stimulation. By estimating parameters of the diffusion model, they established that the effect of MT stimulation on RTs is best explained by change in sensory evidence, while the effect of LIP stimulation is best explained by change in the integrated evidence. This analysis supports the model in which the activity of LIP neurons represent the difference between integrated information from MT neurons.

The models shown in Fig. 18.1 can also fit the time-courses of neural activity in LIP neurons during the choice process. In particular, due to the inhibitory connections, the models can describe the decrease in firing rate of LIP neurons representing the losing alternative [22, 24, 33]. However, the current neurophysiological data do not allow us to distinguish which type of inhibition (feed-forward or mutual) is present in the integrator networks. We will come back to this question in Section 18.5.1.

18.3 Model of decision making in the cortico-basal-ganglia circuit

This section reviews a theory [34] suggesting that the cortico-basal-ganglia circuit implements a generalization of SPRT to choice between multiple alternatives, the Multihypothesis SPRT (MSPRT) [35]. This section follows the same organization as Section 18.2.

18.3.1 Neurobiology of decision processes in the basal ganglia

The basal ganglia are a set of nuclei connected with one another, cortex and subcortical regions. Redgrave et al. [36] have proposed that the basal ganglia act as a central switch resolving competition between cortical regions vying for behavioral expression. In the default state the output nuclei of the basal ganglia send tonic inhibition to the thalamus and brain stem, and thus block execution of any actions [37, 38]. To execute a movement, the firing rates of the corresponding neurons in the output nuclei need to decrease [37, 38], thus releasing the corresponding motor plan from inhibition.

Another property of basal ganglia organization relevant to the model is that within each nucleus different neurons are selective for different body parts [39, 40]. On the basis of this observation it has been proposed that basal ganglia are organized into channels corresponding to individual body parts that traverse all nuclei [41]. Two sample channels are shown in two colors in Fig. 18.2.

Furthermore, the connectivity between nuclei is usually within channels [41]; for example, neurons in the motor cortex selective for the right hand project to the neurons in striatum selective for the right hand, etc. The only exception is the subthalamic nucleus (STN) where neurons project more diffusely across channels [43, 44] (Fig. 18.2).

Figure 18.2 shows a subset of basal ganglia connectivity required for optimal decision making [34]. The figure does not show “the indirect pathway” between the striatum and the output nuclei via the globus pallidus (GP). It has been suggested that this pathway is involved in learning from punishments [45] (see Chapter 19) and will not be discussed in this chapter for simplicity, but it can be included in the model and it continues to implement MSPRT [34].

18.3.2 Optimal stopping criterion for multiple alternatives

Let us first describe MSPRT intuitively (a more formal description will follow). As mentioned in Section 18.2.2, to optimize performance, the choice criterion should take into account the amount of conflict in the evidence. Consequently, in the MSPRT, a choice is made when the integrated evidence for one of the alternatives exceeds a certain level, but this level is not fixed; rather it is increased when the evidence is more conflicting:

$$y_i > \textit{Threshold} + \textit{Conflict}, \quad (18.6)$$

where *Threshold* is a fixed value. Subtracting *Conflict* from both sides, the criterion for making a choice can be written as:

$$y_i - \textit{Conflict} > \textit{Threshold}. \quad (18.7)$$

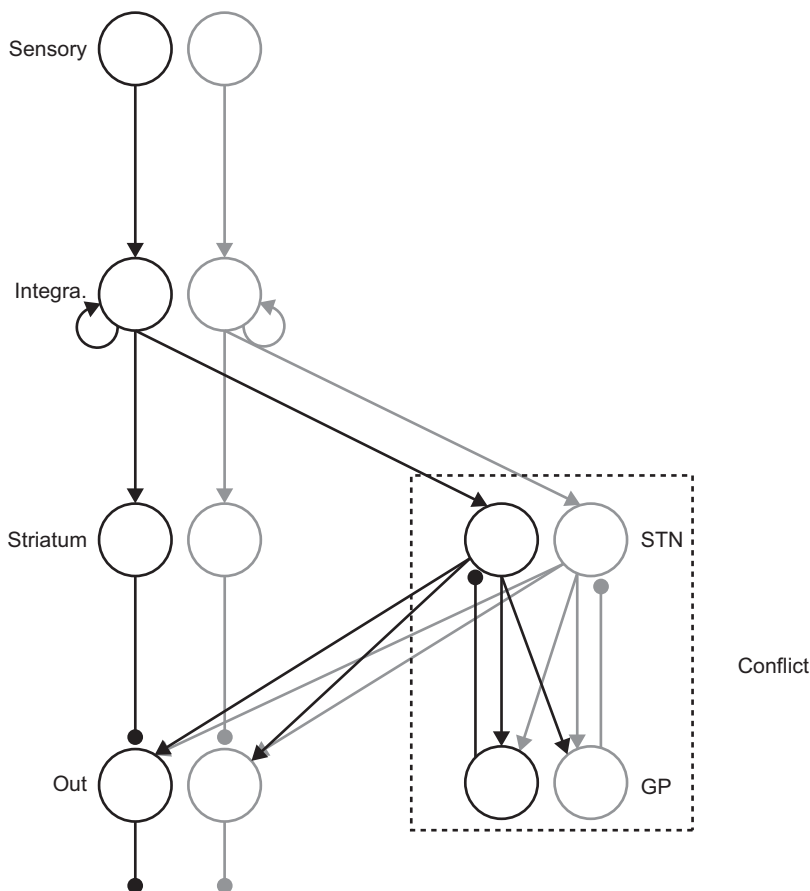


Figure 18.2 A subset of the cortico-basal-ganglia circuit required to implement MSPRT. Connectivity between areas based on Gurney et al. [42]. Pairs of circles correspond to brain areas: Sensory – sensory cortex encoding relevant aspects of stimuli (e.g., MT in motion discrimination task), Integra. – cortical region integrating sensory evidence (e.g., LIP in tasks with saccadic response), STN – subthalamic nucleus, GP – globus pallidus (or its homolog GPe in primates), Out – output nuclei: substantia nigra pars reticulata and entopeduncular nucleus (or its homolog GPi in primates). Arrows denote excitatory connections; lines ended with circles denote inhibitory connections. Black and gray pathways correspond to two sample channels.

In order to implement the MSPRT, *Conflict* needs to take a particular form:

$$Conflict = \log \sum_{j=1}^N \exp y_j, \tag{18.8}$$

where N is the number of alternative choices (we use the italicized word *Conflict* to refer to a term defined in Eq. 18.8, and the non-italicized word “conflict” to refer to a general situation where sensory evidence supports more than one alternative). Equation 18.8 is

said to express conflict because it involves summation of evidence across alternatives, but it also includes particular nonlinearities, and to understand where they come from, we need to describe MSPRT more formally.

MSPRT is a statistical test between N hypotheses. We can generalize the hypotheses described in Eq. 18.2, so that the hypothesis H_i states that the sensory evidence supporting alternative i has the highest mean:

$$H_i: \mu_i = \mu^+, \mu_{j \neq i} = \mu^- . \tag{18.9}$$

In the MSPRT, at each time t , and for each alternative i , one computes the probability of alternative i being correct given all sensory evidence observed so far; let us denote this probability by P_i :

$$P_i = P(H_i | x(1..t)) . \tag{18.10}$$

If for any alternative P_i exceeds a fixed threshold, the choice is made in favor of the corresponding alternative; otherwise the sampling continues. Appendix 18C shows that the logarithm of P_i is equal to:

$$\log P_i = y_i - \log \sum_{j=1}^N \exp y_j . \tag{18.11}$$

Equation 18.11 includes two terms: the integrated evidence y_i and *Conflict* defined in Eq. 18.8. Thus MSPRT is equivalent to making a choice when the difference between y_i and *Conflict* exceeds a threshold.

MSPRT has similar optimality property as SPRT; namely, it minimizes the decision time for any required accuracy (which has been shown analytically for accuracies close to 100% [47], and simulations demonstrate that MSPRT achieves shorter decision times than simpler models also for lower accuracies [48]). Thus, given the discussion in Section 18.2.2 and Appendix 18B, MSPRT also maximizes the reward rate.

18.3.3 Neural implementation

Bogacz and Gurney [34] proposed that the circuit of Fig. 18.2 computes the expression $y_i - \text{Conflict}$ (required for MSPRT). In their model, y_i are computed in cortical integrators accumulating input from sensory neurons (Fig. 18.2), and the *Conflict* is computed by STN and GP (we will describe how it can be computed later). The output nuclei receive inhibition from cortical integrators via the striatum and excitation from STN, thus they compute in the model:

$$OUT_i = -y_i + \text{Conflict} = -(y_i - \text{Conflict}) . \tag{18.12}$$

Thus the output nuclei in the model represent the negative of the expression that is compared against a threshold in MSPRT (cf. Eq. 18.7). Therefore, in the model the choice is made when the activity of any output channel OUT_i decreases below the threshold, in agreement with selection by disinhibition reviewed in Section 18.3.1.

Frank [49] has also proposed that the conflict is computed in STN (see Chapter 19). STN has a suitable anatomical location to modulate choice process according to the

conflict, as it can effectively inhibit all motor programs by its diffuse projections to output nuclei [50]. Studies of patients who receive deep brain stimulation to STN support the idea that STN computes conflict. They suggest that disrupting computations in STN makes patients unable to prevent premature responding in high conflict choices [51, 52] (see Chapter 19 for details).

Bogacz and Gurney [34] showed that input sent by STN to the output nuclei can be proportional to the particular form of *Conflict* defined in Eq. 18.8, if the neurons in STN and GP had the following relationships between their input and the firing rate:

$$STN = \exp(input), \quad (18.13)$$

$$GP = input - \log(input). \quad (18.14)$$

Let us analyze how each term in Eq. 18.8 can be computed. For ease of reading let us restate Eq. 18.8:

$$Conflict = \log \sum_{j=1}^N \exp y_j. \quad (18.8)$$

Starting from the right end of Eq. 18.8, the integrated evidence y_j is provided to the STN in the model by a direct connection from cortical integrators (Fig. 18.2). The exponentiation is performed by the STN neurons (cf. Eq. 18.13). The summation across channels is achieved due to the diffuse projections from the STN. In the model, each channel in the output nuclei receives input from all channels in the STN; hence, the input to the output neurons is proportional to the sum of activity in STN channels. The only non-intuitive element of the computation of Eq. 18.8 is the logarithm—it is achieved by the interactions between STN and GP, as shown in Appendix 18D.

18.3.4 Comparison with experimental data

The input-output relationships of Eqs 18.13 and 18.14 form predictions of the model. The first prediction says that the firing rate of STN neurons should be proportional to the exponent of their input. Figure 18.3 shows the firing rate as a function of injected current for seven STN neurons, for which this relationship has been studied precisely [53, 54]. Solid lines show fits of the exponential function to firing rates below 135 Hz, suggesting that up to approximately 135 Hz the STN neurons have an input-output relationship that is very close to exponential.

For the entire range in Fig. 18.3, the input-output relationship seems to be sigmoidal, often used in neural network models, and one could ask how unique the exponential relationship discussed here is, as every sigmoidal curve has a segment with an exponential increase. However, note that typical neurons have much lower maximum firing rate, so even if they had an exponential segment, it would apply to a much smaller range of firing rates. By contrast, the STN neurons have exponential input-output relationship for up to ~ 135 Hz, and this is approximately the operating range of these neurons in humans during choice tasks [55].

The model also predicts that GP neurons should have approximately linear input-output relationship (because the linear term in Eq. 18.14 dominates for higher *input*).

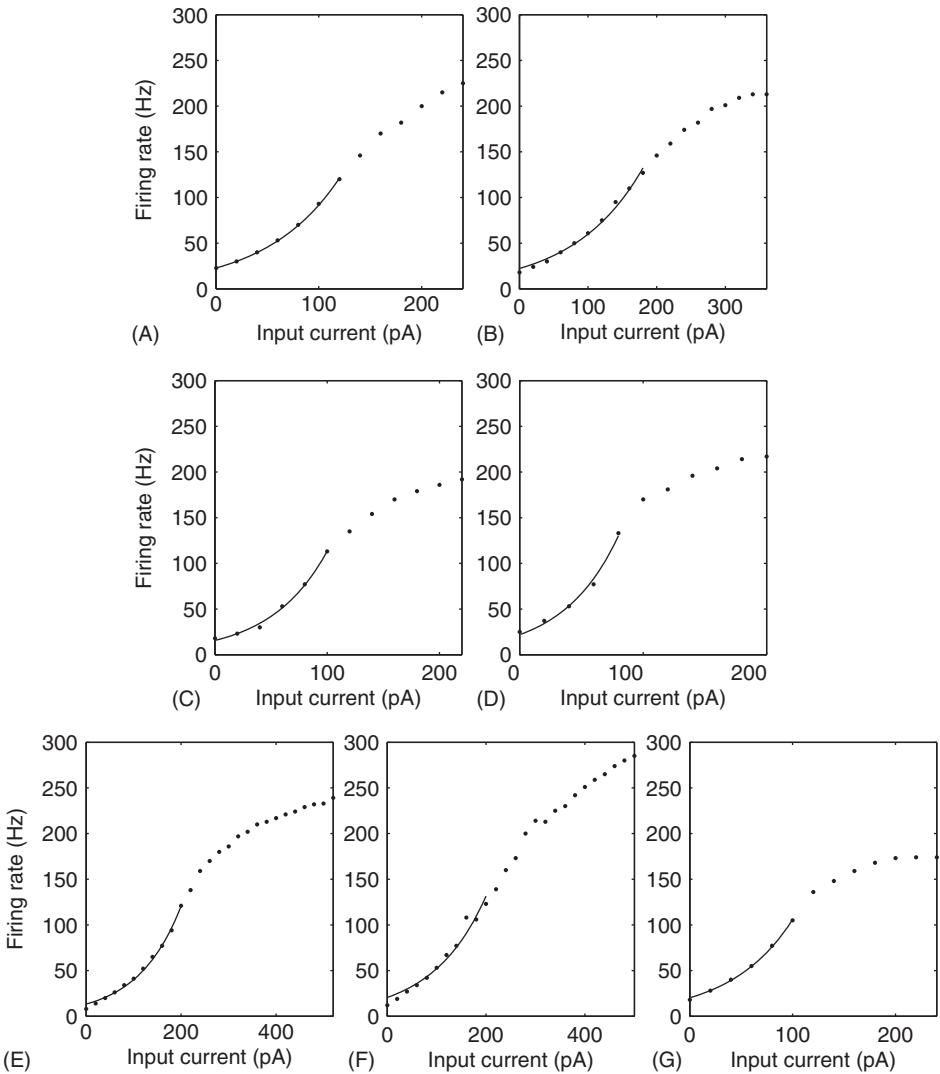


Figure 18.3 Firing rates f of STN neurons as a function of input current I . Panels A–D re-plot data on the firing rate of STN neurons presented in Hallworth et al. [53] in Figures 4b, 4f, 12d, and 13d, respectively (control condition). Panels E–G re-plot the data from STN presented in Wilson et al. [54] in Figures 1c, 2c, and 2f, respectively (control condition). Lines show best fit of the function $f = a \exp(b I)$ to the points with firing rates below 135 Hz.

Neurophysiological studies suggest that there are three distinct subpopulations of neurons in GP, and the one which contributes most to the population firing rate has indeed the linear input-output relationship [56, 57].

The model is also consistent with behavioral data from choice tasks. First, for two alternatives the model produces the same behavior as the diffusion model, because then MSPRT reduces to SPRT, thus the model fits all behavioral data that the diffusion model

can describe (see Section 18.2.4). Second, the model reproduces the Hick's law, stating that the decision time is proportional to the logarithm of the number of alternatives [34]. Third, Norris [58] has shown that MSPRT describes patterns of RTs during word recognition (which can be interpreted as a perceptual decision with the number of alternatives equal to the number of words known by participants).

18.4 Basal ganglia and cortical integration

This section discusses the relationship between the models presented in the two previous sections, and how the basal ganglia may modulate cortical integration allowing cortical integrators to represent probabilities of corresponding alternatives being correct.

18.4.1 Inhibition between cortical integrators

In the model shown in Fig. 18.2, a cortical integrator accumulates sensory evidence only for the corresponding alternative. This simple model is inconsistent with the observation that integrators representing the losing alternative decrease their firing rate during a choice process [2]; to account for this observation the inhibitory connections need to be introduced, as shown in Fig. 18.1.

Nevertheless, if the simple cortical integration in the model of the cortico-basal-ganglia circuit is replaced by any of the models in Fig. 18.1, the circuit still implements MSPRT [34], as we shall now explain. Note that in the biologically more realistic models of Fig. 18.1C and 18.1D, both integrators receive the same inhibition from a pool of inhibitory neurons. The basal ganglia model has a surprising property that if the same inhibition, *inh*, is applied to all integrators, the activity of output nuclei does not change. Intuitively, this property is desirable for the basal ganglia network because, as we discussed in Section 18.2.2, the optimal choice criterion should be based on differences between integrated evidence for the alternatives, so increasing or decreasing the activity of all integrators should not affect the optimal choice criterion. This property is shown formally in Appendix 18E.

18.4.2 Representing probabilities in integrators' firing rate

Recently a modified version of the cortico-basal-ganglia model has been proposed [46] which only differs from the one described in Section 18.8 in that, the integration is performed via cortico-basal-ganglia-thalamic loops as shown in Fig. 18.4. In this modified model the circuit of STN and GP (represented by a small circle in Fig. 18.4) provides the indirect inhibition to the integrators which is the same for all integrators. Due to the property described in the previous paragraph, the activity of output nuclei in the model of Fig. 18.4 is exactly the same as in the model of Fig. 18.2, thus the former also implements MSPRT [46].

Let us now investigate the properties of the cortical integrators predicted by the model of Fig. 18.4. Let us recall that in the model the activities of output nuclei are equal to (from Eqs 18.8, 18.11, and 18.12):

$$OUT_i = -\log P_i \quad (18.15)$$

(recall that P_i denotes the probability of alternative i being correct given the sensory evidence). The integrators in Fig. 18.4 receive input from the sensory neurons and effective




Figure 18.4 Connectivity in the model of integration in cortico-basal-ganglia thalamic loops. Notation as in Figs. 18.1 and 18.2. The dotted rectangle indicates the basal ganglia which are modeled in exactly the same way as in Fig. 18.2, but here for simplicity not all nuclei are shown: small circle represents STN and GP, and cortical integrators send effective inhibition to the output nuclei via the striatum not shown for simplicity.

inhibition from the output nuclei. If there is no sensory evidence coming, the activity of the integrators corresponding to alternative i is proportional to $\log P_i$ (from Eq. 18.15). Thus in this model the cortical integrators represent $\log P_i$ in their firing rate, they update $\log P_i$ according to new sensory evidence, and the basal ganglia continuously renormalizes the cortical activity such that P_i sum up to 1.

18.4.3 Incorporating prior probabilities

Let us now investigate how one can incorporate into the model of Fig. 18.4 the prior probabilities, that is, expectations about correct alternative prior to stimulus onset (e.g., that may arise in perceptual choice task when one stimulus is presented on a greater fraction of trials than others). Before the sensory evidence is provided, P_i are equal to the prior probabilities. Hence, given the discussion in the previous paragraph to utilize the




Figure 18.5 Firing rates f of superior colliculus neurons before stimulus onset as a function of the number of alternatives N . The circles correspond to the experimental data taken from Figure 4 of Basso and Wurtz [62] averaged across two time periods (the position of each dot was computed from the heights of bars in the original figure as $[\text{black bar} + 2 \times \text{gray bar}]/3$, because the gray bars were showing the average firing rate on two times longer intervals than the black bars). The solid line shows the best fitting logarithmic function $f = a \log N + b$.

prior information in the decision process, the initial values of the integrators (just before stimulus onset) should be equal to the logarithm of the prior probabilities.

18.4.4 Comparison with experimental data

The theory reviewed in this section predicts that the activity of integrator neurons should be proportional to $\log P_i$. Two studies [59, 60] have shown directly that activity of LIP neurons is modulated by P_i prior to and during the decision process respectively. But their published results do not allow us to distinguish if the LIP neurons represent $\log P_i$ or some other function of P_i . Nevertheless, there are two less direct studies supporting the hypothesis that the activity of cortical integrators prior to stimulus onset is proportional to the logarithms of the prior probabilities.

First, Carpenter and Williams [61] have also proposed that the starting point of integration is proportional to the logarithm of the prior probability, on the basis of careful analysis of RTs. In their experiment, participants were required to make an eye movement to a dot appearing on the screen. Carpenter and Williams [61] observed that the median RT to a dot appearing in a particular location was proportional to:

$$\text{median RT} \sim -\log \text{Prior}, \quad (18.16)$$

where *Prior* was computed as the fraction of trials within a block in which the dot appears in this location. Now notice that if there were no noise in sensory evidence, then RT would always be constant and proportional to the distance between the starting point (firing rate of integrators at stimulus onset) and the threshold (firing rate at the moment of decision). If the noise is present, RT differs between trials, but the median RT is equal

to the distance between the starting point and the threshold (this relation is exact in a particular model used by Carpenter and Williams [61] and approximate in other models):

$$\text{median RT} \sim \text{Threshold} - \text{starting point} \quad (18.17)$$

Putting Eqs 18.16 and 18.17 together implies that the starting point is proportional to $\log \text{Prior}$.

Second, Basso and Wurtz [62] recorded neural activity in the superior colliculus (a sub-cortical structure receiving input from cortical integrators) in a task in which the number of alternative choices N differed between blocks. Note that in this task $\text{Prior} = 1/N$, and thus $\log \text{Prior} = -\log N$. Figure 18.5 shows that the firing rate before stimulus onset in their task was approximately proportional to $-\log N$, and hence to $\log \text{Prior}$.

18.5 Discussion

In this chapter we reviewed theories proposing that cortical networks and cortico-basal-ganglia circuit implement optimal statistical tests for decision making on the basis of noisy information. The predictions of these theories have been shown to be consistent with both behavioral and neurophysiological data. We have also reviewed a theory suggesting that the basal ganglia modulate cortical integrators allowing them to represent probabilities of corresponding alternatives being correct. In this section we discuss open questions.

18.5.1 Open questions

Let us discuss four open questions relating to the theories reviewed in this chapter.

1. Is the integration of sensory evidence supported by feedback loops within the cortex (Fig. 18.1) or by the cortico-basal-ganglia-thalamic loops (Fig. 18.4)? These possibilities could be distinguished by deactivation of striatal neurons (in the relevant channels) during a decision task. Then, if information is integrated via the cortico-basal-ganglia-thalamic loops, the gradual increasing firing rates of cortical integrators should no longer be observed.
2. What type of inhibition is present between cortical integrators? In this chapter we reviewed three possible pathways by which cortical integrators can compete and inhibit one another, shown in Figs. 18.1C, 18.1D, and 18.4. It should be possible to distinguish between these models on the basis of future neurophysiological studies. For example, in tasks in which the amount of evidence for the two alternatives can be varied independently, the feed-forward inhibition model predicts that the activity of the integrators should only depend on the difference between sensory evidence, while the mutual-inhibition model predicts that it should also depend on the total input to the integrators [15].
3. What else is represented by the cortical integrators? The theories reviewed in this chapter propose that the integrators encode integrated evidence or the probability of the corresponding alternative being correct. However, it is known that the firing rate of the cortical integrators is also modulated by other factors, for example, motor preparation or general desirability of alternatives [63]. It would be interesting to investigate how these other factors can further optimize decision making.

4. What is the relationship between decision making and reinforcement learning in the basal ganglia? The basal ganglia are also strongly involved in reinforcement learning (see Chapter 19). It would be interesting to integrate reinforcement learning and decision-making theories [46]. Furthermore, it could be interesting to model the role of dopamine during decision process, and how it modulates information processing in the basal ganglia during choice.

18A Appendix A: Diffusion model implements SPRT

This appendix describes the relationship between SPRT and the diffusion model. First note that if we assume that $x_i(t)$ are sampled independently, then the likelihood of the sequence of samples is equal to the product of the likelihoods of individual samples:

$$P(x(1..t)|H_1) = \prod_{\tau=1}^t P(x(\tau)|H_1), \tag{18A.1}$$

where $x(\tau)$ denotes a pair: $x_1(\tau), x_2(\tau)$. Substituting Eq. 18A.1 into Eq. 18.3 and taking the logarithm, we obtain:

$$\log R = \sum_{\tau=1}^t \log \frac{P(x(\tau)|H_1)}{P(x(\tau)|H_2)}. \tag{18A.2}$$

We will now illustrate how $\log R$ changes during the choice process. First note that according to Eq. 18A.2, at each time step t , a term is added to $\log R$ which only depends on $x(t)$. This term will be positive if $x_1(t) > x_2(t)$, because then $x(t)$ will be more likely given hypothesis H_1 (stating that $\mu_1 > \mu_2$) than given hypothesis H_2 , thus $P(x(t)|H_1) > P(x(t)|H_2)$. Hence in each time step, $\log R$ increases if $x_1(t) > x_2(t)$, and analogously decreases if $x_1(t) < x_2(t)$.

Let us now derive the likelihood ratio given in Eq. 18.3 for the hypotheses of Eq. 18.2. Let us start with the numerator of Eq. 18.3. Hypothesis H_1 states that $x_1(t)$ come from a normal distribution with mean μ^+ , while $x_2(t)$ come from a normal distribution with mean μ^- (see Eq. 18.2). Thus denoting the probability density of normal distribution with mean μ and standard deviation σ by $f_{\mu,\sigma}$ Eq. 18A.1 becomes:

$$P(x(1..t)|H_1) = \prod_{\tau=1}^t f_{\mu^+,\sigma}(x_1(\tau))f_{\mu^-,\sigma}(x_2(\tau)). \tag{18A.3}$$

Using the equation for the normal probability density function, we obtain:

$$P(x(1..t)|H_1) = \prod_{\tau=1}^t \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{(x_1(\tau) - \mu^+)^2}{2\sigma^2}\right) \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{(x_2(\tau) - \mu^-)^2}{2\sigma^2}\right). \tag{18A.4}$$

Performing analogous calculation for the denominator of Eq. 18.3 we obtain very similar expression as in Eq. 18A.4 but with swapped μ^+ and μ^- . Hence, when we write the ratio of the numerator and the denominator, many terms will cancel. First, the constants in front of the exponents will cancel and we obtain:

$$R = \prod_{\tau=1}^t \frac{\exp\left(-\frac{(x_1(\tau) - \mu^+)^2}{2\sigma^2}\right) \exp\left(-\frac{(x_2(\tau) - \mu^-)^2}{2\sigma^2}\right)}{\exp\left(-\frac{(x_1(\tau) - \mu^-)^2}{2\sigma^2}\right) \exp\left(-\frac{(x_2(\tau) - \mu^+)^2}{2\sigma^2}\right)}. \quad (18A.5)$$

Second, if we expand the squares inside the exponents, and split the exponents using $\exp(a + b) = \exp(a)\exp(b)$, most of the terms will cancel, and we obtain:

$$R = \prod_{\tau=1}^t \frac{\exp\left(\frac{\mu^+ x_1(\tau)}{\sigma^2}\right) \exp\left(\frac{\mu^- x_2(\tau)}{\sigma^2}\right)}{\exp\left(\frac{\mu^- x_1(\tau)}{\sigma^2}\right) \exp\left(\frac{\mu^+ x_2(\tau)}{\sigma^2}\right)}. \quad (18A.6)$$

Now if we join the terms including $x_1(\tau)$ using $\exp(a)/\exp(b) = \exp(a - b)$, and do the same for the terms including $x_2(\tau)$ we obtain:

$$R = \prod_{\tau=1}^t \exp(gx_1(\tau))\exp(-gx_2(\tau)), \quad \text{where } g = \frac{\mu^+ - \mu^-}{\sigma^2}. \quad (18A.7)$$

If we put the product inside the exponent, it becomes a summation, and using Eq. 18.1 we obtain:

$$R = \frac{\exp(gy_1)}{\exp(gy_2)}. \quad (18A.8)$$

In summary, there were many cancellations during the derivation, and comparing Eqs 18.3 and 18A.8 reveals that the only term which remains from $P(x(1..t)|H_i)$ is the exponent of the evidence integrated in favor of alternative i , and the exponentiation comes from the exponent in the normal probability density function. Taking the logarithm of Eq. 18A.8 gives Eq. 18.4.

18B Appendix B: Diffusion model maximizes reward rate

This appendix shows that the diffusion model with optimally chosen threshold achieves higher or equal reward rate than any other model of decision making. As mentioned in Section 18.2.2, decision-making models exhibit a speed-accuracy tradeoff (controlled by the decision threshold). For a given task, let us denote the average RT of the diffusion

model for a given accuracy by $RT_d(AC)$, and the RT of another model by $RT_a(AC)$. The optimality of SPRT implies that for any accuracy level AC :

$$RT_d(AC) \leq RT_a(AC). \tag{18B.1}$$

Similarly, let us denote the reward rates of the diffusion and the other model for a given level of accuracy by $RR_d(AC)$ and $RR_a(AC)$ respectively. Let us denote the accuracy level that maximizes the reward rate of the diffusion model by AC_d , so from this definition for any accuracy level AC :

$$RR_d(AC) \leq RR_a(AC_d). \tag{18B.2}$$

Similarly, let us denote the accuracy maximizing the reward rate of the other model by AC_a . We can now derive the following relationship between the maximum possible reward rates achieved by the two models:

$$RR_a(AC_a) = \frac{AC_a}{RT_a(AC_a) + D} \leq \frac{AC_a}{RT_d(AC_a) + D} = RR_d(AC_a) \leq RR_d(AC_d). \tag{18B.3}$$

In the above derivation, the first inequality comes from Inequality Eq. 18B.1, and the second inequality from Inequality Eq. 18B.2. Thus Inequality Eq. 18B.3 implies that the highest reward rate possible for the diffusion model, $RR_d(AC_d)$, is higher or equal than the highest possible reward rate in any other model, $RR_a(AC_a)$.

18C Appendix C: MSPRT

This appendix shows that the logarithm of the probability P_i defined in Eq. 18.10 is expressed by Eq. 18.11. We can compute P_i from the Bayes theorem:

$$P_i \equiv P(H_i | x(1..t)) = \frac{P(x(1..t) | H_i)P(H_i)}{P(x(1..t))}. \tag{18C.1}$$

As usual in statistical testing, we assume that one of the hypotheses H_i is correct, thus the probability of sensory evidence $P(x(1..t))$ is equal to the average probability of sensory evidence given the hypotheses (weighted by the prior probabilities of the hypotheses):

$$P_i = \frac{P(x(1..t) | H_i)P(H_i)}{\sum_{j=1}^N P(x(1..t) | H_j)P(H_j)}. \tag{18C.2}$$

Let us for simplicity assume that we do not have any prior knowledge favoring any of the alternatives, so the prior probabilities are equal (Section 18.4.3 shows how the prior

probabilities can be incorporated). Then the prior probabilities in Eq. 18C.2 cancel and it becomes:

$$P_i = \frac{P(x(1..t) | H_i)}{\sum_{j=1}^N P(x(1..t) | H_j)}. \quad (18C.3)$$

Now, recall that in Appendix 18A we have already evaluated a very similar ratio for a similar set of hypotheses. Following calculations like those in Appendix 18A, analogous cancellations happen, and Eq. 18C.3 becomes (cf. Eq. 18A.8):

$$P_i = \frac{\exp(gy_i)}{\sum_{j=1}^N \exp(gy_j)}, \quad (18C.4)$$

where g is a constant defined in Eq. 18A.7. As we observed in Appendix 18A, the exponentiation of integrated evidence comes from the exponents present in the normal probability density function. If we take the logarithm of Eq. 18C.4, then the division becomes a subtraction:

$$\log P_i = \log \exp(gy_i) - \log \sum_{j=1}^N \exp(gy_j). \quad (18C.5)$$

Equation 18C.5 involves two terms: In the first term the logarithm cancels the exponentiation so it becomes the integrated evidence (scaled by constant g), while the second is the *Conflict* in evidence (scaled by g). For simplicity, we ignore the constant g , because its precise value has been shown to be of little importance for performance [34], and then Eq. 18C.5 becomes Eq. 18.11.

18D Appendix D: STN and GP can compute the optimal conflict

In this appendix, we show that if STN and GP have input-output relationship given by Eqs 18.13 and 18.14, then the sum of the activities of STN channels is equal to *Conflict*. Let us denote the activities of channel i in STN and GP by STN_i and GP_i respectively. The STN receives input from cortical integrators and inhibition from GP, hence

$$STN_i = \exp(y_i - GP_i). \quad (18D.1)$$

Let us denote the sum of activity of all STN channels by Σ :

$$\Sigma = \sum_{j=1}^N STN_j. \quad (18D.2)$$

In the model of Fig. 18.2, GP receives input from all STN channels, hence

$$GP_i = \Sigma - \log \Sigma. \quad (18D.3)$$

Substituting Eq. 18D.3 into Eq. 18D.1 gives

$$STN_i = \exp(y_i - \Sigma + \log \Sigma). \quad (18D.4)$$

Using the property of exponentiation $e^{a+b} = e^a e^b$ we obtain:

$$STN_i = \exp(y_i) \exp(-\Sigma) \Sigma. \quad (18D.5)$$

Summing over i and using Eq. 18D.2 we obtain:

$$\Sigma = \sum_{i=1}^N \exp(y_i) \exp(-\Sigma) \Sigma. \quad (18D.6)$$

Taking the logarithm of Eq. 18D.6 we get:

$$\log \Sigma = \log \sum_{i=1}^N \exp y_i - \Sigma + \log \Sigma. \quad (18D.7)$$

$\log \Sigma$ cancels on both sides in Eq. 18D.7, and moving Σ on the left side we see that the sum of activities of all STN channels is equal to *Conflict*:

$$\Sigma = \log \sum_{i=1}^N \exp y_i. \quad (18D.8)$$

18E Appendix E: Basal ganglia model is unaffected by inhibition of integrators

This appendix shows that if the same inhibition is applied to all cortical integrators projecting to the basal ganglia, the activity of the output nuclei in the model does not change. To see this property, we subtract *inh* from y_i in the first line of Eq. 18E.1 (cf. Eqs 18.8 and 18.12) and observe that *inh* cancels out (last equality in Eq. 18E.1) and the activity of the output nuclei does not change.

$$\begin{aligned} OUT_i &= -(y_i - inh) + \log \left(\sum_{j=1}^N \exp(y_j - inh) \right) = \\ &= -y_i + inh + \log \left(\left(\sum_{j=1}^N \exp y_j \right) \exp(-inh) \right) = \\ &= -y_i + inh + \log \left(\sum_{j=1}^N \exp y_j \right) + \log \exp(-inh) = -y_i + \log \left(\sum_{j=1}^N \exp y_j \right). \end{aligned} \quad (18E.1)$$

Acknowledgement

This work was supported by EPSRC grants EP/C514416/1 and EP/C516303/1. The author thanks Simon Chiu for reading the previous version of the manuscript and for his very useful comments.

References

- [1] K.H. Britten, M.N. Shadlen, W.T. Newsome, J.A. Movshon, Responses of neurons in macaque MT to stochastic motion signals, *Vis. Neurosci.* 10 (1993) 1157–1169.
- [2] M.N. Shadlen, W.T. Newsome, Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey, *J. Neurophysiol.* 86 (2001) 1916–1936.
- [3] J.D. Roitman, M.N. Shadlen, Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task, *J. Neurosci.* 22 (2002) 9475–9489.
- [4] J. Ditterich, M. Mazurek, M.N. Shadlen, Microstimulation of visual cortex affect the speed of perceptual decisions, *Nat. Neurosci.* 6 (2003) 891–898.
- [5] J.D. Schall, Neural basis of deciding, choosing and acting, *Nat. Rev. Neurosci.* 2 (2001) 33–42.
- [6] J.I. Gold, M.N. Shadlen, The neural basis of decision making, *Annu. Rev. Neurosci.* 30 (2007) 535–574.
- [7] P. Cisek, J.F. Kalaska, Neural correlates of reaching decisions in dorsal premotor cortex: specification of multiple direction choices and final selection of action, *Neuron* 45 (2005) 801–814.
- [8] H.R. Heekeren, S. Marrett, P.A. Bandettini, L.G. Ungerleider, A general mechanism for perceptual decision-making in the human brain, *Nature* 431 (2004) 859–862.
- [9] K.G. Thompson, D.P. Hanes, N.P. Bichot, J.D. Schall, Perceptual and motor processing stages identified in the activity of macaque frontal eye field neurons during visual search, *J. Neurophysiol.* 76 (1996) 4040–4055.
- [10] H.R. Heekeren, S. Marrett, L.G. Ungerleider, The neural systems that mediate human perceptual decision making, *Nat. Rev. Neurosci.* 9 (2008) 467–479.
- [11] D. Vickers, Evidence for an accumulator model of psychophysical discrimination, *Ergonomics* 13 (1970) 37–58.
- [12] D.R.J. Laming, *Information Theory of Choice Reaction Time*, Wiley, New York, 1968.
- [13] R. Ratcliff, A theory of memory retrieval, *Psychol. Rev.* 83 (1978) 59–108.
- [14] M. Stone, Models for choice reaction time, *Psychometrika* 25 (1960) 251–260.
- [15] R. Bogacz, Optimal decision-making theories: linking neurobiology with behaviour, *Trends Cogn. Sci.* 11 (2007) 118–125.
- [16] J.I. Gold, M.N. Shadlen, Neural computations that underlie decisions about sensory stimuli, *Trends Cogn. Sci.* 5 (2001) 10–16.
- [17] J.I. Gold, M.N. Shadlen, Banburismus and the brain: decoding the relationship between sensory stimuli, decisions, and reward, *Neuron* 36 (2002) 299–308.
- [18] A. Wald, *Sequential Analysis*, Wiley, New York, 1947.
- [19] A. Wald, J. Wolfowitz, Optimum character of the sequential probability ratio test, *Ann. Math. Stat.* 19 (1948) 326–339.
- [20] R. Bogacz, E. Brown, J. Moehlis, P. Holmes, J.D. Cohen, The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced choice tasks, *Psychol. Rev.* 113 (2006) 700–765.
- [21] D.P. Hanes, J.D. Schall, Neural control of voluntary movement initiation, *Science* 274 (1996) 427–430.

- [22] M.E. Mazurek, J.D. Roitman, J. Ditterich, M.N. Shadlen, A role for neural integrators in perceptual decision making, *Cereb. Cortex* 13 (2003) 1257–1269.
- [23] M. Usher, J.L. McClelland, The time course of perceptual choice: the leaky, competing accumulator model, *Psychol. Rev.* 108 (2001) 550–592.
- [24] X.J. Wang, Probabilistic decision making by slow reverberation in cortical circuits, *Neuron* 36 (2002) 955–968.
- [25] K.-F. Wong, X.-J. Wang, A recurrent network mechanism of time integration in perceptual decisions, *J. Neurosci.* 26 (2006) 1314–1328.
- [26] R. Ratcliff, R. Gomez, G. McKoon, A diffusion model account of the lexical decision task, *Psychol. Rev.* 111 (2004) 159–182.
- [27] R. Ratcliff, J.N. Rouder, A diffusion model account of masking in two-choice letter identification, *J. Exp. Psychol. Hum. Percept. Perform.* 26 (2000) 127–140.
- [28] R. Ratcliff, A. Thapar, G. McKoon, A diffusion model analysis of the effects of aging on brightness discrimination, *Percept. Psychophys.* 65 (2003) 523–535.
- [29] R. Ratcliff, A. Cherian, M. Segraves, A comparison of macaques behavior and superior colliculus neuronal activity to predictions from models of two-choice decisions, *J. Neurophysiol.* 90 (2003) 1392–1407.
- [30] R. Ratcliff, P.L. Smith, Comparison of sequential sampling models for two-choice reaction time, *Psychol. Rev.* 111 (2004) 333–367.
- [31] P.L. Smith, R. Ratcliff, Psychology and neurobiology of simple decisions, *Trends Neurosci.* 27 (2004) 161–168.
- [32] T.D. Hanks, J. Ditterich, M.N. Shadlen, Microstimulation of macaque area LIP affects decision-making in a motion discrimination task, *Nat. Neurosci.* 9 (2006) 682–689.
- [33] J. Ditterich, Stochastic models of decisions about motion direction: behavior and physiology, *Neural Netw.* 19 (2006) 981–1012.
- [34] R. Bogacz, K. Gurney, The basal ganglia and cortex implement optimal decision making between alternative actions, *Neural Comput.* 19 (2007) 442–477.
- [35] C.W. Baum, V.V. Veeravalli, A sequential procedure for multihypothesis testing, *IEEE Trans. Inf. Theory* 40 (1994) 1996–2007.
- [36] P. Redgrave, T.J. Prescott, K. Gurney, The basal ganglia: a vertebrate solution to the selection problem? *Neuroscience* 89 (1999) 1009–1023.
- [37] G. Chevalier, S. Vacher, J.M. Deniau, M. Desban, Disinhibition as a basic process in the expression of striatal functions. I. The striato-nigral influence on tecto-spinal/tecto-diencephalic neurons, *Brain Res.* 334 (1985) 215–226.
- [38] J.M. Deniau, G. Chevalier, Disinhibition as a basic process in the expression of striatal functions. II. The striato-nigral influence on thalamocortical cells of the ventromedial thalamic nucleus, *Brain Res.* 334 (1985) 227–233.
- [39] A.P. Georgopoulos, M.R. DeLong, M.D. Crutcher, Relations between parameters of step-tracking movements and single cell discharge in the globus pallidus and subthalamic nucleus of the behaving monkey, *J. Neurosci.* 3 (1983) 1586–1598.
- [40] E. Gerardin, S. Lehericy, J.B. Pochon, S.T. Montcel, J.F. Mangin, F. Poupon, Y. Agid, D. Le Bihan, C. Marsault, Foot, hand, face and eye representation in human striatum, *Cereb. Cortex* 13 (2003) 162–169.
- [41] G.E. Alexander, M.R. DeLong, P.L. Strick, Parallel organization of functionally segregated circuits linking basal ganglia and cortex, *Annu. Rev. Neurosci.* 9 (1986) 357–381.
- [42] K. Gurney, T.J. Prescott, P. Redgrave, A computational model of action selection in the basal ganglia. I. A new functional anatomy, *Biol. Cybern.* 84 (2001) 401–410.
- [43] A. Parent, L.N. Hazrati, Anatomical aspects of information processing in primate basal ganglia, *Trends Neurosci.* 16 (1993) 111–116.

- [44] A. Parent, Y. Smith, Organization of efferent projections of the subthalamic nucleus in the squirrel monkey as revealed by retrograde labeling methods, *Brain Res.* 436 (1987) 296–310.
- [45] M.J. Frank, L.C. Seeberger, R.C. O'Reilly, By carrot or by stick: cognitive reinforcement learning in parkinsonism, *Science* 306 (2004) 1940–1943.
- [46] R. Bogacz, optimal decision-making in cortico-basal-ganglia-thalamic loops, forum of European Neuroscience, Geneva, Abstracts, 126.2 (2008).
- [47] V.P. Dragalin, A.G. Tertakovskiy, V.V. Veeravalli, Multihypothesis sequential probability ratio tests – part I: asymptotic optimality, *IEEE Trans. Inf. Theory* 45 (1999) 2448–2461.
- [48] T. McMillen, P. Holmes, The dynamics of choice among multiple alternatives, *J. Math. Psychol.* 50 (2006) 30–57.
- [49] M.J. Frank, Hold your horses: a dynamic computational role for the subthalamic nucleus in decision making, *Neural Netw.* 19 (2006) 1120–1136.
- [50] J.W. Mink, The basal ganglia: focused selection and inhibition of competing motor programs, *Prog. Neurobiol.* 50 (1996) 381–425.
- [51] M.J. Frank, J. Samanta, A.A. Moustafa, S.J. Sherman, Hold your horses: impulsivity, deep brain stimulation, and medication in parkinsonism, *Science* 318 (2007) 1309–1312.
- [52] M. Jahanshahi, C.M. Ardouin, R.G. Brown, J.C. Rothwell, J. Obeso, A. Albanese, M.C. Rodriguez-Oroz, E. Moro, A.L. Benabid, P. Pollak, P. Limousin-Dowsey, The impact of deep brain stimulation on executive function in Parkinson's disease, *Brain* 123 (Pt 6) (2000) 1142–1154.
- [53] N.E. Hallworth, C.J. Wilson, M.D. Bevan, Apamin-sensitive small conductance calcium-activated potassium channels, through their selective coupling to voltage-gated calcium channels, are critical determinants of the precision, pace, and pattern of action potential generation in rat subthalamic nucleus neurons in vitro, *J. Neurosci.* 23 (2003) 7525–7542.
- [54] C.J. Wilson, A. Weyrick, D. Terman, N.E. Hallworth, M.D. Bevan, A model of reverse spike frequency adaptation and repetitive firing of subthalamic nucleus neurons, *J. Neurophysiol.* 91 (2004) 1963–1980.
- [55] Z.M. Williams, J.S. Neimat, G.R. Cosgrove, E.N. Eskandar, Timing and direction selectivity of subthalamic and pallidal neurons in patients with Parkinson disease, *Exp. Brain Res.* 162 (2005) 407–416.
- [56] H. Kita, S.T. Kitai, Intracellular study of rat globus pallidus neurons: membrane properties and responses to neostriatal, subthalamic and nigral stimulation, *Brain Res.* 564 (1991) 296–305.
- [57] A. Nambu, R. Llinas, Electrophysiology of globus pallidus neurons in vitro, *J. Neurophysiol.* 72 (1994) 1127–1139.
- [58] D. Norris, The Bayesian reader: explaining word recognition as an optimal Bayesian decision process, *Psychol. Rev.* 113 (2006) 327–357.
- [59] M.L. Platt, P.W. Glimcher, Neural correlates of decision variables in parietal cortex, *Nature* 400 (1999) 233–238.
- [60] T. Yang, M.N. Shadlen, Probabilistic reasoning by neurons, *Nature* 447 (2007) 1075–1080.
- [61] R.H. Carpenter, M.L. Williams, Neural computation of log likelihood in control of saccadic eye movements, *Nature* 377 (1995) 59–62.
- [62] M.A. Basso, R.H. Wurtz, Modulation of neuronal activity in superior colliculus by changes in target probability, *J. Neurosci.* 18 (1998) 7519–7534.
- [63] M.C. Dorris, P.W. Glimcher, Activity in posterior parietal cortex is correlated with the relative subjective desirability of action, *Neuron* 44 (2004) 365–378.