NeuroImage 53 (2010) 694–706

Contents lists available at ScienceDirect

NeuroImage

journal homepage: www.elsevier.com/locate/ynimg

Choice, difficulty, and confidence in the brain

Edmund T. Rolls [a,*], Fabian Grabenhorst [b], Gustavo Deco [c]

[a] Oxford Centre for Computational Neuroscience, Oxford, England
[b] Department of Experimental Psychology, University of Oxford, South Parks Road, Oxford OX1 3UD, England
[c] Institució Catalana de Recerca i Estudis Avançats (ICREA), Universitat Pompeu Fabra, Computational Neuroscience, Barcelona, Spain

ARTICLE INFO

Article history:
Received 13 January 2010
Revised 23 April 2010
Accepted 25 June 2010
Available online 6 July 2010

Keywords:
Decision-making
Confidence
Attractor network
Task difficulty
fMRI
Computational neuroscience
Emotion
Temperature
Odor
Medial prefrontal cortex area 10
Reaction time
Reward value

ABSTRACT

To provide a neurobiological basis for understanding decision-making and decision confidence, we describe and analyze a neuronal spiking attractor-based model of decision-making that makes predictions about synaptic and neuronal activity, the fMRI BOLD response, and behavioral choice as a function of the easiness of the decision, and thus decision confidence. The spiking network model predicts probabilistic decision-making with faster and larger neuronal responses on easy versus difficult choices, that is as the discriminability ΔI between the choices increases, and these and the synaptic currents in turn predict larger BOLD responses as the discriminability increases. Confidence, which increases with discriminability, thus emerges from the firing rates of the decision-making neurons in the choice attractor network. In two fMRI studies, we confirm these predictions by showing that brain areas such as medial prefrontal cortex area 10 implicated in choice decision-making between pleasant stimuli have BOLD activations linearly related to the easiness of both olfactory and warm pleasantness choices. Further, this signature is not found in orbitofrontal cortex areas that represent on a continuous scale the value of the stimuli, but are not implicated in the choice itself. This provides a unifying and fundamental approach to decision-making and decision confidence, and to how spiking-related noise in the brain affects choice, confidence, synaptic and neuronal activity, and fMRI signals.

Introduction

Decision-making, and our confidence in the decisions we make, are important areas in neuroscience in the new field of neuroeconomics (Behrens et al., 2007; Deco and Rolls, 2006; Glimcher, 2003; Hampton and O'Doherty, 2007; Heekeren et al., 2004; Kepecs et al., 2008; Kiani and Shadlen, 2009; Kim and Shadlen, 1999; Knutson et al., 2007; Marsh et al., 2007; Rolls, 2008; Romo et al., 2004; Shadlen and Newsome, 1996, 2001; Sugrue et al., 2005). Critical issues are to understand the neural processes that underlie decision-making and decision confidence, and to establish the bases for neural signatures of these processes. Primate single neuron recordings (Kim and Shadlen, 1999) show that neuronal responses in a motion decision-making task occur earlier on easy vs. difficult trials in a decision-related brain region, the dorsolateral prefrontal cortex. In the human dorsolateral prefrontal cortex, higher fMRI BOLD signals can be observed on easy trials vs. difficult trials in a similar decision-making task, and this has been proposed as a signature of decision-making (Heekeren et al., 2004; Heekeren et al., 2008). However, exactly what neuronal mechanisms might underlie such decision-making, and whether this is a reliable signature, are not yet clear. The accumulator, counter, or race models of decision-making in which the evidence for different choices accumulates (Carpenter and Williams, 1995; Ratcliff et al., 1999; Smith and Ratcliff, 2004; Usher and McClelland, 2001; Vickers, 1979; Vickers and Packer, 1982) do not describe a neurobiological mechanism, and do not make specific predictions about how neuronal firing rates, synaptic currents, or BOLD signals will alter as decisions become easier and decision confidence increases.

Here we capitalize on recent advances in theoretical understanding of how choice decisions are made using an integrate-and-fire attractor network that makes probabilistic decisions from the spontaneous firing state into one of two or more high firing rate stable attractor states each implemented by a set of coupled neurons that receives the inputs for one of the decisions, and where the choice made is probabilistic because of the noise contributed to by the random spiking times of the neurons for a given firing rate (Deco and Rolls, 2006; Deco et al., 2009; Rolls, 2008; Rolls and Deco, 2010; Wang, 2002, 2008). Such attractor networks are implemented by excitatory connections between cortical pyramidal cells, and provide a neural architecture not only for decision-making but also for short-term memory (Goldman-Rakic, 1995; Rolls, 2008) and memory recall (Rolls, 2008). We show that the model predicts higher firing rates, synaptic currents, and fMRI BOLD (functional magnetic resonance neuroimaging blood oxygenation level dependent) signals

* Corresponding author.
E-mail address: Edmund.Rolls@oxcns.org (E.T. Rolls).
URL: http://www.oxcns.org (E.T. Rolls).

1053-8119/$ – see front matter © 2010 Elsevier Inc. All rights reserved.
doi:10.1016/j.neuroimage.2010.06.073

on easy trials vs. difficult trials, and in fact that these increase monotonically (approximately linearly) with ΔI, the difference between the stimuli. Further, it is well established that subjective decision confidence increases with discriminability, ΔI (Jonsson et al., 2005; Vickers, 1979; Vickers and Packer, 1982), and thus the degree of confidence in a decision emerges from the model as being reflected by the firing rates of the neurons involved in the decision-making. Consistently, in rats too the probability that a trial will be aborted reflecting low decision confidence also increases with ΔI on error trials (Kepecs et al., 2008). Moreover, neuronal data consistent with the predictions of the model about decision confidence have been recorded by Kiani and Shadlen (2009) in the posterior parietal cortex during a perceptual decision-making task, and neurons in the rat orbitofrontal cortex that respond on error trials with firing rates that increase with ΔI, that is with confidence in the error (Kepecs et al., 2008), may also be consistent (Insabato et al., 2010; Rolls and Deco, 2010). We test these predictions of the attractor network theory of decision-making in two fMRI investigations of decision-making about the reward value and subjective pleasantness of thermal and olfactory stimuli.

Methods

Modelling investigations

The theoretical framework of the model used here was introduced by Wang (2002) and developed further (Deco and Rolls, 2006; Deco et al., 2009; Deco et al., 2007; Marti et al., 2008; Rolls and Deco, 2010; Wang, 2008), and the results described here apply generically to integrate-and-fire attractor network models of decision-making. In this framework, we model probabilistic decision-making by a network of interacting neurons organized into a discrete set of populations, as depicted in Fig. 1. Populations or pools of neurons are defined as groups of excitatory or inhibitory neurons sharing the same inputs and connectivities. The network contains N_E (excitatory) pyramidal cells and N_I inhibitory interneurons. In our simulations, we use $N_E = 400$ and $N_I = 100$, consistent with the neurophysiologically observed proportion of 80% pyramidal cells versus 20% interneurons (Abeles, 1991; Rolls and Deco, 2002). The neurons are fully connected (with synaptic strengths as specified below). In the model, the specific populations D1 (for decision 1) and D2 encode the categorical result of the choice, between the two stimuli that activate each of these populations. Each specific population of excitatory cells contains rN_E neurons (in our simulations $r = 0.1$). In addition there is one non-specific population, named "Non-specific," which groups all other excitatory neurons in the modelled brain area not involved in the present tasks, and one inhibitory population, named "Inhibitory," grouping the local inhibitory neurons in the modelled brain area. The latter population regulates the overall activity and implements competition in the network by spreading a global inhibition signal.

Because we are mainly interested in the nonstationary probabilistic behaviour of the network, the proper level of description at the microscopic level is captured by the spiking and synaptic dynamics of one-compartment *Integrate-and-Fire* (IF) neuron models (Deco and Rolls, 2005). An IF neuron integrates the afferent current generated by the incoming spikes, and fires when the depolarization of the cell membrane crosses a threshold. At this level of detail the model allows the use of realistic biophysical time constants, latencies and conductances to model the synaptic current, which in turn allows a thorough study of the realistic time scales and firing rates involved in the time evolution of the neural activity. Consequently, the simulated neuronal dynamics, that putatively underlie cognitive processes, can be quantitatively compared with experimental data. For this reason, it is very useful to include a thorough description of the different time constants of the synaptic activity. The IF neurons are modelled as having three types of receptor mediating the synaptic currents flowing into them: AMPA, NMDA (both activated by glutamate), and GABA receptors. The excitatory recurrent post-synaptic currents (EPSCs) are considered to be mediated by AMPA (fast) and NMDA (slow) receptors; external EPSCs imposed onto the network from outside are modelled as being driven only by AMPA receptors. Inhibitory post-synaptic currents (IPSCs) to both excitatory and inhibitory neurons are mediated by GABA receptors. The details of the mathematical formulation are summarized in previous papers (Brunel and Wang, 2001; Deco and Rolls, 2006), and are provided in the Supplementary Material.

We modulate the conductance values for the synapses between pairs of neurons by connection weights, which can deviate from their default value 1. The structure and function of the network are achieved by differentially setting the weights within and between populations of neurons. The labelling of the weights is defined in Fig. 1. We assume that the connections are already formed, by for example earlier self-organization mechanisms, as if they were established by Hebbian learning, i.e. the coupling will be strong if the pair of neurons have correlated activity (i.e. covarying firing rates), and weak if they are activated in an uncorrelated way. As a consequence of this, neurons within a specific excitatory population (D1 and D2) are mutually coupled with a strong synaptic weight w_+, set to 2.1 for the simulations described here. Furthermore, the populations encoding for these two decisions are likely to have anti-correlated activity in this behavioral context, resulting in weaker than average connections between the two different populations. Consequently, we choose a weaker value $w_- = 1 - r(w_+ - 1)/(1 - r)$, so that the overall recurrent excitatory synaptic drive in the spontaneous state remains constant as w_+ is varied (Brunel and Wang, 2001). Neurons in the inhibitory population are mutually connected with an intermediate weight $w = 1$. They are also connected with all excitatory neurons in the same layer with the same intermediate weight, which for excitatory-to-inhibitory connections is $w = 1$, and for inhibitory-to-excitatory connections is denoted by a weight w_I. Neurons in a specific excitatory population are connected to neurons in the nonselective population in the same layer with a feedforward synaptic weight $w = 1$ and a feedback synaptic connection of weight w_-.

Each individual population is driven by two different kinds of input. First, all neurons in the model network receive spontaneous background activity from outside the module through $N_{ext} = 800$ external excitatory connections. Each synaptic connection carries a Poisson spike train at a spontaneous rate of 3 Hz, which is a typical spontaneous firing rate value observed in the cerebral cortex. This results in a background external input summed over all 800 synapses of 2400 Hz for each neuron. Second, the neurons in the two specific populations additionally receive added firing to the external inputs that encode the evidence for the decision to be made. When stimulating, the rate of the Poisson train to the neurons of the specific population D1 is increased by an extra value of λ_1, and to population D2 by λ_2, as these encode the two stimuli to be compared.

The simulations were run for 2 s of spontaneous activity, and then for a further 2 s while the stimuli were being applied. During the spontaneous period, the stimuli applied to D1 and D2 (and to all the other neurons in the network) had a value of 3 Hz. (This 3 Hz is the firing rate being applied by Poisson spikes to all 800 external synaptic inputs of each neuron in D1, so the total synaptic bombardment on each neuron is 2400 spikes/s.) During the decision period, the mean input to D1 and D2 was increased to 3.04 Hz per synapse (an extra 32 Hz per neuron). For $\Delta I = 0$, we added 32 extra Hz to the spontaneous and applied this to each neuron of both D1 and D2. For $\Delta I = 16$, 32 + 8 Hz was the extra applied to D1 and corresponds to λ_1 in Fig. 1, and 32–8 Hz was the extra applied to D2, etc. The firing rates, and the absolute value of the sum of the synaptic currents (AMPA, NMDA, and GABA, defined in the Supplementary Material), in all four populations of neurons were saved every 50 ms for later analysis. The criterion for which population won, that is for which decision was taken, was a mean rate for the last second of the simulation that was

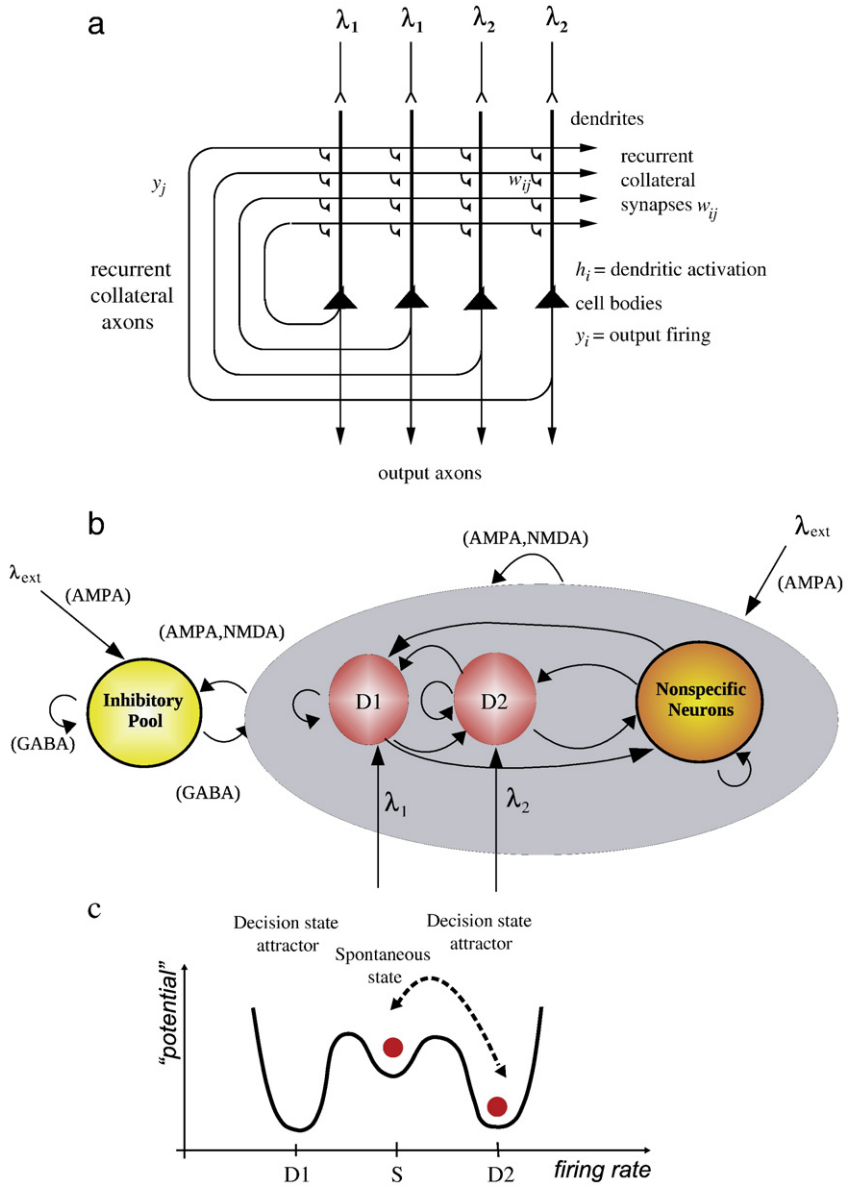

Fig. 1. (a). Attractor or autoassociation single network architecture for decision-making. The evidence for decision 1 is applied via the λ_1 inputs, and for decision 2 via the λ_2 inputs. The synaptic weights w_{ij} have been associatively modified during training in the presence of λ_1 and at a different time of λ_2. When λ_1 and λ_2 are applied, each attractor competes through the inhibitory interneurons (not shown), until one wins the competition, and the network falls into one of the high firing rate attractors that represents the decision. The noise in the network caused by the random spiking of the neurons means that on some trials, for given inputs, the neurons in the decision 1 (D1) attractor are more likely to win, and on other trials the neurons in the decision 2 (D2) attractor are more likely to win. This makes the decision-making probabilistic, for, as shown in (c), the noise influences when the system will jump out of the spontaneous firing stable (low energy) state S, and whether it jumps into the high firing state for decision 1 (D1) or decision 2 (D2). (b) The architecture of the integrate-and-fire network used to model decision-making (see text). (c) A multistable "effective energy landscape" for decision-making with stable states shown as low "potential" basins. Even when the inputs are being applied to the network, the spontaneous firing rate state is stable, and noise provokes transitions into the high firing rate decision attractor state D1 or D2 (see Rolls and Deco 2010).

10 Hz greater than that of the other population (this is in the context that the spontaneous rate was several spikes/s, and that the winning population typically had a mean firing rate of 35–40 Hz, as will be shown). The latency of a decision was measured by the time of the first 50 ms bin of three consecutive bins at which the mean rate of one population was more than 25 Hz higher than that of the other population. The haemodynamic signal associated with the decision was calculated by convolving the neuronal activity or the synaptic currents of the neurons with the haemodynamic response function used by SPM5 (Statistical Parametric Mapping, Wellcome Trust Centre for Neuroimaging, London) (as this was the function also used in the analyses by SPM of the experimental fMRI data). For the convolution, the predecision period of spontaneous activity was padded out so that it lasted for 30 s, and after the 2-s period of decision-making activity, each trial was padded out for a further 18 s with spontaneous activity, so that the effects found in a 2 s period of decision-making could be measured against a steady background. The predicted BOLD signals are shown with $t=0$ corresponding to the time when the decision stimuli were turned on, just as in the analyses of the experimental fMRI data, to enable direct comparisons. Qualitatively similar results, though with a somewhat slower timecourse, were obtained with a biophysically based balloon model for the generation of fMRI BOLD signals (Stephan et al., 2007). We note that synaptic currents are not necessarily linearly related to the firing rate of neurons due for

example to neuronal firing threshold non-linearities, and address this further in the Supplementary Material).

The parameters for the synaptic weights and input currents were chosen using the mean-field equivalent of this network (Brunel and Wang, 2001; Deco and Rolls, 2006) so that in the absence of noise when the input stimuli are being applied there were three stable states, the spontaneous firing rate state (with a mean firing for the pyramidal cells of approximately 3 spikes/s), and two high firing rate attractor states (with a mean firing for the pyramidal cells of approximately 40 spikes/s), with one neuronal population (D1) representing decision 1, and the other population (D2) decision 2. In particular, w_+ was set to 2.1.

The particular model chosen using mean field analysis (Brunel and Wang, 2001; Deco and Rolls, 2006; Rolls and Deco, 2010) had three stable states in the absence of the spiking noise, a spontaneous state and one for each of two decisions, as this model has many interesting decision-making properties, including implementing Weber's law ($\Delta I / I = k$, i.e. the difference of intensity I that can be reliably detected divided by the intensity value has a linear component) (Deco and Rolls, 2006; Deco et al., 2009). With operation in this multistable regime, it is the approximately random spiking times of the neurons (i.e. approximately Poisson firing at a given mean rate) that causes statistical fluctuations that makes the network jump from the spontaneous firing state into one of the high firing rate attractor (decision) states. The randomness of the firing dynamically and probabilistically provokes transitions from the spontaneous firing state to one of the high firing rate attractor basins that represent a decision (see Fig. 1c).

The stability of an attractor is characterized by the average time in which the system stays in the basin of attraction under the influence of noise. The noise provokes transitions to other attractor states. It results from the interplay between the Poissonian character of the spikes and the finite-size effect due to the limited numbers of neurons in the network (Rolls and Deco, 2010).

Olfactory pleasantness decision task

The design is based on the vibrotactile decision-making studies of Romo and colleagues in which vibrotactile stimuli separated by a delay were presented, and a decision had to be made when the second stimulus was presented of whether the second frequency was higher than the first (Romo et al., 2004). In our design the decision was about whether the second odor was more pleasant than the first, or on other trials was more intense than the first. To allow a comparison with trials on which choices between stimuli were not made, there were also trials in which only ratings of the continuous affective value and intensity were made, without a choice between stimuli, as described in detail by Rolls et al. (2010). The odors were delivered through a computer-controlled olfactometer. The pleasant odors were 1 M citral and 4 M vanillin. The unpleasant odors were hexanoic acid (10% v/v) and isovaleric acid (15%).

The experimental protocol consisted of an event-related interleaved design presenting in random permuted sequence the 3 experimental conditions and different pairs of olfactory stimuli for each condition. Each trial started at $t = 0$ s with the first odor being delivered for 2 s accompanied by a visual label stating "Sniff first stimulus." There was then a 6-s period during which clean air was delivered. In this period at $t = 7$ s a visual label was displayed stating either "Decide pleasantness," "Decide intensity" or "Rate stimulus." At $t = 8$ s the second odor was presented for 2 s accompanied by a visual label stating either "Sniff Decide pleasantness," "Sniff Decide intensity" or "Sniff Rate stimulus". There was then a 6-s period during which clean air was delivered. Starting at $t = 16$ s on decision trials, the words "First stimulus" and "Second stimulus" appeared on the screen for 2 s, and in this period the participant had to select which button key response to make (up button or down button) for the decision that had been taken at the time when the second odor was delivered. At $t = 18$ s on decision trials, the word "First stimulus" or "Second stimulus" was then displayed to provide feedback to the participant that their choice was acknowledged. On rating trials, starting at $t = 16$ s the subjective ratings were made. The first rating was for the pleasantness of the second odor on a continuous visual scale from −2 (very unpleasant), through 0 (neutral), to +2 (very pleasant). The second rating was for the intensity of the second odor on a scale from 0 (very weak) to 4 (very intense). The ratings were made with a visual rating scale in which the subject moved the bar to the appropriate point on the scale using a button box. There was 4 s for each rating. Each of the trial types was presented in random permuted sequence 36 times. Which two of the four odors were presented on each trial, and the order in which they were presented, was determined by a random permuted sequence.

ΔI, the difference in pleasantness of the two stimuli between which a decision was being made, was obtained for each trial by the absolute value of the difference in the (average) rated pleasantness of that pair of stimuli for each subject. Thus, two odors of similar pleasantness would have a small ΔI, and two odors of different pleasantness would have a large ΔI. This measure thus reflects the difficulty of the decision, and is independent of whether the second odor happened to be pleasant or unpleasant. This value for ΔI on every trial was used as a parametric regressor in the fMRI analyses, to identify brain regions with activations positively correlated with ΔI, the easiness of the affective decision. (An analogous procedure was used to calculate ΔI for the trials on which the decision was about intensity, but this is not further discussed here as brain areas with significant effects for easy vs. difficult decisions about intensity were not identified).

Temperature pleasantness decision task

Warm and cool thermal stimuli, and mixtures of them, were applied to the hand. In a previous investigation of the same dataset, we compared brain responses when participants were taking decisions about whether they would select a thermal stimulus (yes vs. no), with activations to the same stimuli on different trials when only affective ratings were required and there was no decision about whether the participants would say yes or no to the stimuli if they were available in the future (Grabenhorst et al., 2008). In the present investigation, we analysed data only on decision trials, and the analyses were about how activations with the stimuli were related to the difficulty of the decision, which was not investigated previously. Both the decision and rating trials were identical from the start of each trial at $t = 0$ s until $t = 5$ s when a visual stimulus was shown for 1 s stating "decide" or "rate" the thermal stimulus being applied, and at $t = 6$ s a green cross appeared until $t = 10$ s. On decide trials from $t = 6$ until $t = 10$ s the participants had to decide whether yes or no was the decision on that trial. At $t = 10$ s a visual stimulus with yes above no or vice versa in random order was shown for 2 s, and the participant had to press the upper or lower button on the button box as appropriate to indicate the response. On rating trials from $t = 6$ until $t = 10$ s the participants had to encode the pleasantness and intensity of the thermal stimulus being applied, so that the ratings could be made later. On rating trials at $t = 10$ s the pleasantness rating could be made using the same button box, and then the intensity rating was made. The thermal stimuli were a warm pleasant stimulus (41 °C) applied to the hand ("warm2"), a cool unpleasant stimulus (12 °C) applied to the hand ("cold"), a combined warm and cold stimulus ("warm2 + cold"), and a second combination designed to be less pleasant (39 °C + 12 °C) ("warm1 + cold"), delivered with Peltier devices as described previously (Grabenhorst et al., 2008).

ΔI for the thermal stimuli was the absolute value of the pleasantness rating, based on the concept that it is more difficult to choose whether a stimulus should be repeated in future if it is close to

neutral (0) in rated pleasantness, vs. is rated as being pleasant (with the maximum pleasantness being +2), or as being unpleasant (with the most unpleasant being -2).

fMRI data acquisition

Images were acquired with a 3.0-T VARIAN/SIEMENS whole-body scanner at the Centre for Functional Magnetic Resonance Imaging at Oxford (FMRIB), where 27 T2* weighted EPI coronal slices with in-plane resolution of 3×3 mm and between plane spacing of 4 mm were acquired every 2 seconds (TR = 2), as described previously for these datasets (Grabenhorst et al., 2008; Rolls et al., 2010). Twelve healthy volunteers (7 male and 5 female, mean age 27) participated in the olfactory decision-making study (Rolls et al., 2010), and twelve (6 male and 6 female, mean age 26) in the temperature decision-making study (Grabenhorst et al., 2008). The studies were conducted in accord with the declaration of Helsinki and approved by the Central Oxford Research Ethics Committee, and written informed consent from all subjects was obtained before the experiment.

fMRI analyses

The imaging data were analysed using SPM5 (Statistical Parametric Mapping, Wellcome Institute of Cognitive Neurology, London). Pre-processing of the data used SPM5 realignment, reslicing with sinc interpolation, normalisation to the MNI coordinate system (Montreal Neurological Institute) (Collins et al., 1994), and spatial smoothing with a 6-mm full width at half maximum isotropic Gaussian kernel. Time series non-sphericity at each voxel was estimated and corrected for (Friston et al., 2002), and a high-pass filter with a cut-off period of 128 s was applied. In the single event design, a general linear model was then applied to the time course of activation where the decision or rating period onsets ($t = 8$ s in each trial for the olfactory task and $t = 6$ s for the temperature task) were modelled as single impulse response functions and then convolved with the canonical hemodynamic response function (Friston et al., 1994). Linear contrasts were defined to test specific effects. Given that the fMRI analyses were to test hypotheses generated by our computational model of decision-related activity in the brain, the main test performed was for brain areas where activations were greater on easy compared to difficult trials, and brain areas identified in this way were then compared to brain areas implicated by other analyses in making choices. In this hypothesis-led research, we were able to find activations consistent with our hypotheses, and to show that such areas are implicated by other analyses in making the choices. Time derivatives were included in the basis functions set. Following smoothness estimation (Kiebel et al., 1999), in the first stage of analysis condition-specific experimental effects (parameter estimates, or regression coefficients, pertaining to the height of the canonical HRF) were obtained via the general linear model (GLM) in a voxel-wise manner for each subject. In the second (group random effects) stage, subject-specific linear contrasts of these parameter estimates were entered into a series of one-sample *t*-tests, each constituting a group-level statistical parametric map. We report results for brain regions where there were prior hypotheses about brain areas involved in decision-making or rating affective value (Grabenhorst et al., 2008; Rolls et al., 2010), and applied small volume (false discovery rate) corrections for multiple comparisons (Genovese et al., 2002) with a radius corresponding to the full width at half maximum of the spatial smoothing filter used. Peaks are reported for which $p < 0.05$, though the exact corrected probability values (Worsley et al., 1996) are given in the text.

For the time course plots shown in Fig. 5, we extracted event-related responses from the peak voxel for each subject using the fitted responses provided by SPM, which reflect the hemodynamic response function and are baseline adjusted. These single-subject time courses were then averaged across subjects.

The criteria used to identify regions involved in choice decision-making in earlier investigations with the same data set were that a brain region should show more brain activity with identical stimuli on trials on which a choice decision was being made than when the continuous affective value of the stimuli were being rated, but no choice was being made between stimuli, or about whether the stimulus would be chosen again (Grabenhorst et al., 2008; Rolls and Grabenhorst, 2008; Rolls et al., 2010). For the olfactory task, a contrast of decision vs. rating trials showed activations in the medial prefrontal cortex (medial area 10) at [2 50 −12] $z = 3.78$ $p < 0.001$. For the thermal task, a contrast of decision vs. rating trials showed activations at [6 54 −8] $z = 3.24$ $p = 0.022$.

Results

An attractor network model of decision-making and confidence

The attractor network model of decision-making analyzed is shown in Fig. 1. The implementation of the integrate-and-fire model with AMPA, NMDA, and GABA dynamic synapses is described in Methods and Supplementary Material.

Fig. 2a and e shows the mean firing rates of the two neuronal populations D1 and D2 for two trial types, easy trials ($\Delta I = 160$ Hz) and difficult trials ($\Delta I = 0$) (where ΔI is the difference in spikes/s summed across all synapses to each neuron between the two inputs, λ_1 to population D1, and λ_2 to population D2). The results are shown for correct trials, that is, trials on which the D1 population won the competition and fired with a rate for >10 spikes/s for the last 1000 ms of the simulation runs (this provided a clear criterion for which attractor won the competition as shown by the binary distribution of the rates of the two attractors found over 1000 simulation trials, and as exemplified by the single trials in Fig. 2). Fig. 2b shows the mean firing rates of the four populations of neurons on a difficult trial, and Fig. 2c shows the rastergrams for the same trial. Fig. 2d shows the firing rates on another difficult trial ($\Delta I = 0$) to illustrate the variability shown from trial to trial, with on this trial prolonged competition between the D1 and D2 attractors until the D1 attractor finally won after approximately 1100 ms. Fig. 2f shows firing rate plots for the 4 neuronal populations on an example of a single easy trial ($\Delta I = 160$), Fig. 2g shows the synaptic currents in the four neuronal populations on the same trial, and Fig. 2h shows rastergrams for the same trial.

Three important points are made by the results shown in Fig. 2. First, the network falls into its decision attractor faster on easy trials than on difficult trials. We would accordingly expect reaction times to be shorter on easy than on difficult trials. We might also expect the BOLD signal related to the activity of the network to be higher on easy than on difficult trials because it starts sooner on easy trials. Second, the mean firing rate after the network has settled into the correct decision attractor is higher on easy than on difficult trials. We might therefore expect the BOLD signal related to the activity of the network to be higher on easy than on difficult trials because the maintained activity in the attractor is higher on easy trials. This shows that the exact firing rate in the attractor is a result not only of the internal recurrent collateral effect, but also of the external input to the neurons, which in Fig. 2 is 32 Hz to each neuron (summed across all synapses) of D1 and D2, but in Fig. 2a is increased by 80 Hz to D1, and decreased by 80 Hz to D2 (i.e. the total external input to the network is the same, but $\Delta I = 0$ for Fig. 2a, and $\Delta I = 160$ for Fig. 2b). Third, the variability of the firing rate is high, with the standard deviations of the mean firing rate calculated in 50 ms epochs indicated in order to quantify the variability. The large standard deviations on difficult trials for the first second after the decision cues are applied at $t = 2$ s reflects the fact that on some trials the network has entered an attractor state after 1000 ms, but on other trials it has not yet reached the attractor, although it does so later. This trial by trial variability is

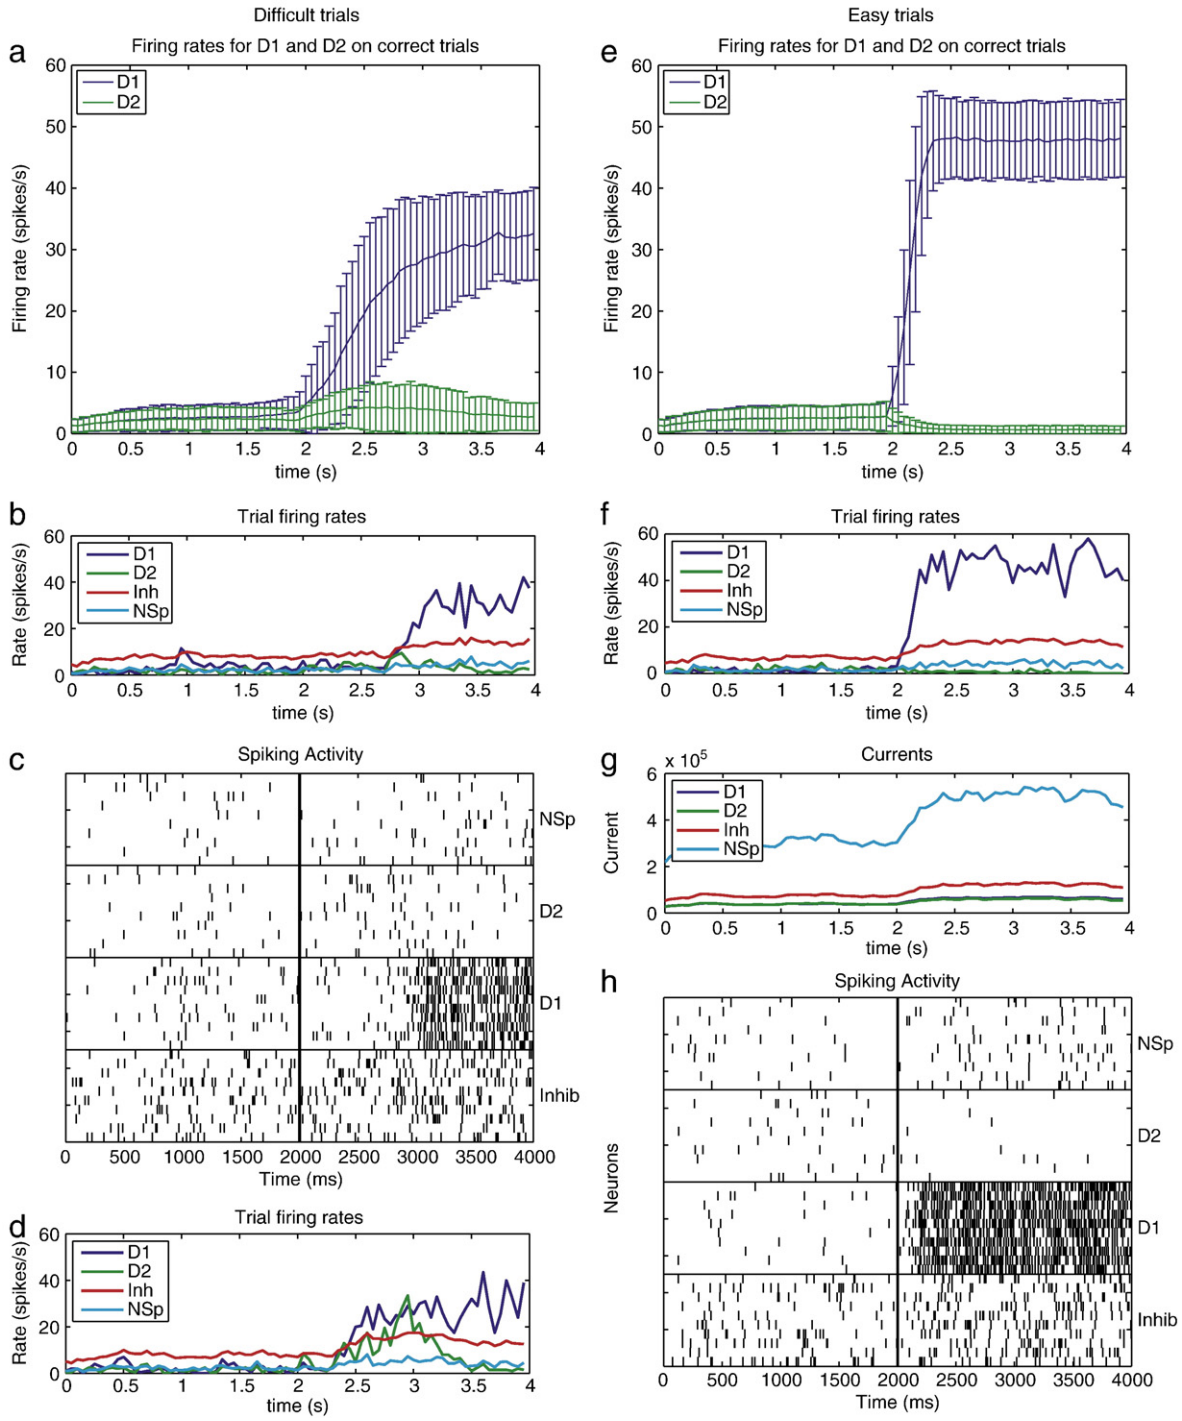

Fig. 2. (a and e) Firing rates (mean ± sd) for difficult ($\Delta I = 0$) and easy ($\Delta I = 160$) trials. The period 0-2 s is the spontaneous firing, and the decision cues were turned on at time = 2 s. The mean was calculated over 1000 trials. D1: firing rate of the D1 population of neurons on correct trials on which the D1 population won. D2: firing rate of the D2 population of neurons on the correct trials on which the D1 population won. A correct trial was one in which in which the mean rate of the D1 attractor averaged >10 spikes/s for the last 1000 ms of the simulation runs (given the attractor nature of the network and the parameters used, the network reached one of the attractors on >90% of the 1000 trials, and this criterion clearly separated these trials, as indicated by the mean rates and standard deviations for the last s of the simulation as shown). (b) The mean firing rates of the four populations of neurons on a difficult trial. Inh is the inhibitory population that uses GABA as a transmitter. NSp is the non-specific population of neurons (see Fig. 1). (c) Rastergrams for the trial shown in b. 10 neurons from each of the four pools of neurons are shown. (d) The firing rates on another difficult trial ($\Delta I = 0$) showing prolonged competition between the D1 and D2 attractors until the D1 attractor finally wins after approximately 1100 ms. (f) Firing rate plots for the 4 neuronal populations on a single easy trial ($\Delta I = 160$). (g) The synaptic currents in the four neuronal populations on the trial shown in f. (h) Rastergrams for the easy trial shown in f and g.

indicated by the firing rates on individual trials and the rastergrams in the lower part of Fig. 2. The effects evident in Fig. 2 are quantified, and elucidated over a range of values for ΔI, next.

Fig. 3a shows the firing rates (mean ± sd) on correct trials when in the D1 attractor as a function of ΔI. $\Delta I = 0$ corresponds to the most difficult decision, and $\Delta I = 160$ corresponds to easy. The firing rates were measured in the last 1 s of firing, i.e. between $t = 3$ and $t = 4$ s. It is clear that the mean firing rate increases monotonically as ΔI increases, and interestingly, the increase is approximately linear (Pearson $r = 0.995$, $p < 10^{-6}$). The higher mean firing rates as ΔI

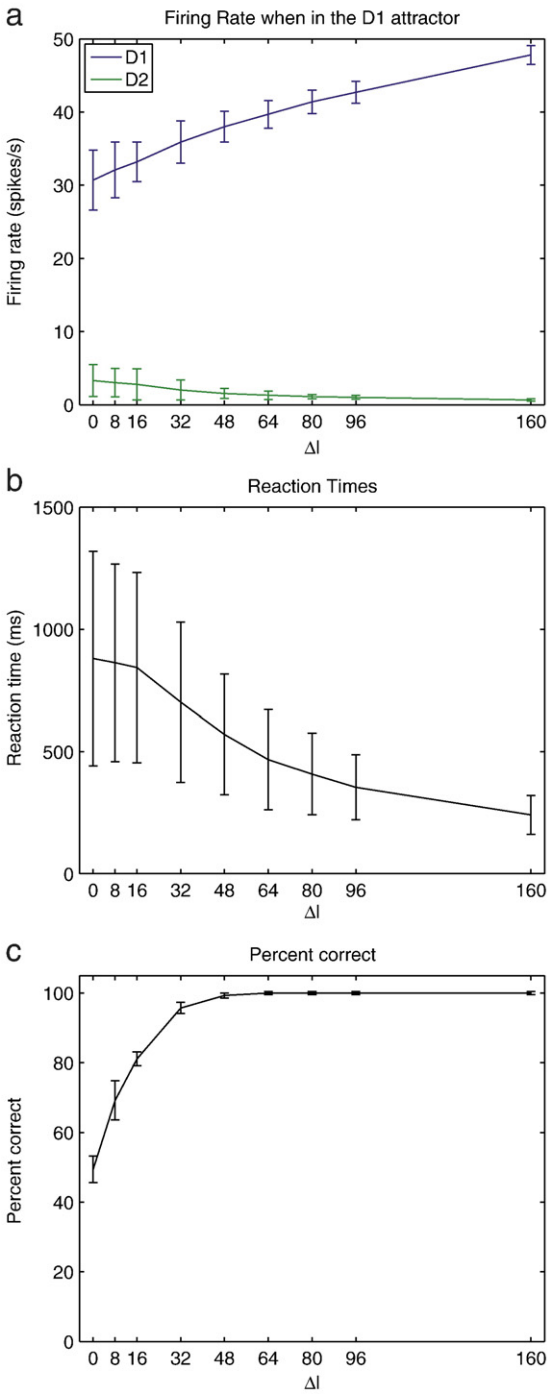

Fig. 3. (a) Firing rates (mean ± sd) when in the D1 attractor as a function of ΔI. $\Delta I = 0$ corresponds to difficult, and $\Delta I = 160$ corresponds to easy. The firing rates for both the winning population D1 and for the losing population D2 on correct trials are shown. All the results are for 1000 simulation trials for each parameter value, and all the results shown are statistically very highly significant. (b) Decision or reaction times (mean ± sd) for the D1 population to win on correct trials as a function of the difference in inputs ΔI to D1 and D2. $\Delta I = 0$ corresponds to difficult, and $\Delta I = 160$ corresponds to easy. (c) Percent correct performance, i.e. the percentage of trials on which the D1 population won, as a function of the difference in inputs ΔI to D1 and D2. The mean was calculated over 1000 trials, and the standard deviation was estimated by the variation in 10 groups each of 100 trials. $\Delta I = 0$ corresponds to difficult, and $\Delta I = 160$ corresponds to easy.

increases are due not only to higher peak firing, but also to the fact that the variability of firing within a trial becomes less as ΔI increases ($r = -0.95$, $p < 10^{-4}$), reflecting the fact that the system is more noisy and unstable with low ΔI, whereas the firing rate in the attractor is maintained more stably with smaller statistical fluctuations against the Poisson effects of the random spike timings at high ΔI (the measure of variation indicated in the figure is the standard deviation, and this is shown throughout unless otherwise stated to quantify the degree of variation, which is a fundamental aspect of the operation of these neuronal decision-making networks). The increase of the firing rate when in the D1 attractor (blue line) as ΔI increases thus reflects the confidence of the decision which increases with ΔI (Jonsson et al., 2005; Vickers, 1979; Vickers and Packer, 1982), and, as will be shown next in Fig. 3b, the performance as shown by the percentage of correct choices. The firing rate of the losing attractor (D2, green line) decreases as ΔI increases, due to feedback inhibition from the winning D1 attractor, and thus the difference in the firing rates of the two attractors also reflects well the decision confidence.

The time for the network to reach the correct D1 attractor, i.e. the decision or reaction time of the network, is shown as a function of ΔI in Fig. 3b (mean ± sd). Interestingly, the reaction time continues to decrease ($r = -0.95$, $p < 10^{-4}$) over a wide range of ΔI, even when as shown in Fig. 3c the network is starting to perform at 100% correct. The decreasing reaction time as ΔI increases is attributable to the altered "effective energy landscape" (Rolls and Deco, 2010) (see Fig. 1c): a larger input to D1 tends to produce occasionally higher firing rates, and these statistically are more likely to induce a significant depression in the landscape towards which the network flows sooner than with low ΔI. Correspondingly, the variability (quantified by the standard deviation) of the reaction times is greatest at low ΔI, and decreases as ΔI increases ($r = -0.95$, $p < 10^{-4}$). This variability would not be found with a deterministic system (i.e. the standard deviations would be 0 throughout, and such systems include those investigated with mean-field analyses), and is entirely due to the random statistical fluctuations caused by the random spiking of the neurons in the integrate-and-fire network.

At $\Delta I = 0$, there is no influence on the network to fall more into attractor D1 representing decision 1 than attractor D2 representing decision 2, and its decisions are at chance, with approximately 50% of decisions being for D1. As ΔI increases, the proportion of trials on which D1 is reached increases. The relation between ΔI and percentage correct is shown in Fig. 3c. Interestingly, the performance becomes 100% correct with $\Delta I = 64$, whereas as shown in Fig. 3a and b the firing rates while in the D1 attractor (and therefore potentially the BOLD signal), continue to increase as ΔI increases further, and the reaction times continue to decrease as ΔI increases further. It is a clear prediction for neurophysiological and behavioral measures that the firing rates with decisions made by this attractor process continue to increase as ΔI is increased beyond the level for very good performance as indicated by the percentage of correct decisions, and the neuronal and behavioral reaction times continue to decrease as ΔI is increased beyond the level for very good performance. Fig. 3c also shows that the variability in the percentage correct (in this case measured over blocks of 100 trials) is large with $\Delta I = 0$, and decreases as ΔI increases. This in consistent with unbiased effects of the noise producing very variable effects in the energy landscape at $\Delta I = 0$, but in the external inputs biasing the energy landscape more and more as ΔI increases, so that the flow is much more likely to be towards the D1 attractor.

Predictions of fMRI BOLD signals from the model

We now show how this model makes predictions for the fMRI BOLD signals that would occur in brain areas in which decision-making processing of the type described is taking place. The BOLD signals were predicted from the firing rates of the neurons in the network (or from the synaptic currents flowing in the neurons as described later) by convolving the neuronal activity with the haemodynamic response function in a realistic period, the 2 s after the decision cues are applied. This is a reasonable period to take, as decisions will be taken within this time, and the attractor state may

not necessarily be maintained for longer than this. The attractor states might be maintained for longer if the response that can be made is not known until later, as in some fMRI tasks with delays including those described here, and then the effects described might be expected to be larger, given the mean firing rate effects shown in Fig. 2. Indeed, it is an advantage of the type of model described here that it is a short-term memory attractor network that can maintain its firing whether or not the decision cues remain, as this continuing firing enables the decision state to be maintained until a response based on it can be guided and made. As shown in Fig. 4a, the predicted fMRI response is larger for easy ($\Delta I = 160$) than for difficult trials ($\Delta I = 0$), with intermediate trials ($\Delta I = 80$) producing an intermediate fMRI response. The difference in the peak response for $\Delta I = 0$ and $\Delta I = 160$ is highly significant ($p << 0.001$). Importantly, the BOLD response is inherently variable from brain regions associated with this type of decision-making process, and this is nothing to do with noise in measuring the response with a scanner. If the system were deterministic, the standard deviations, shown as a measure of the variability in Fig. 4a, would be 0. It is the statistical fluctuations caused by the noisy (random) spike timings of the neurons that account for the variability in the BOLD signals in Fig. 4a. Interestingly, the variability is larger on the difficult trials ($\Delta I = 0$) than on the easy trials ($\Delta I = 160$), as shown in Fig. 4a, and indeed this also can be taken as an indicator that attractor decision-making processes of the type described here are taking place in a brain region.

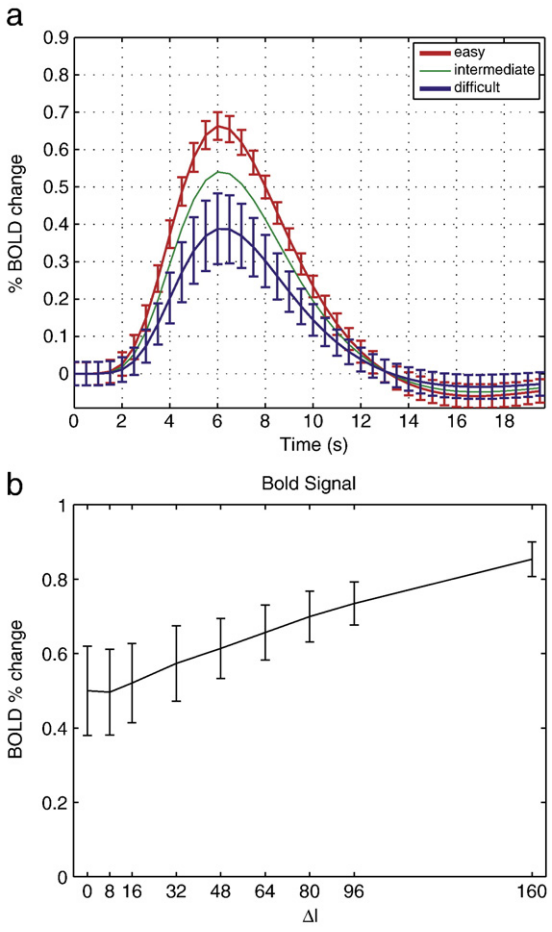

Fig. 4. (a) The percentage change in the BOLD signal on easy trials ($\Delta I = 160$), on intermediate trials ($\Delta I = 80$), and on difficult trials ($\Delta I = 0$). The mean ± sd are shown for the easy and difficult trials. The percentage change in the BOLD signal was calculated from the firing rates of the D1 and D2 populations, and analogous effects were found with calculation from the synaptic currents averaged for example across all 4 populations of neurons. (b) The percentage change in the BOLD signal (peak mean ± sd) as a function of ΔI. $\Delta I = 0$ corresponds to difficult, and $\Delta I = 160$ corresponds to easy. The percent change was measured as the change from the level of activity in the simulation period 1–2 s, before the decision cues were applied at $t = 2$ s, and was calculated from the firing rates of the neurons in the D1 and D2 populations. The BOLD percent change scaling is arbitrary, and is set so that the lowest value is 0.5%.

Fig. 4b shows that the percentage change in the BOLD signal (peak mean ± sd) predicted from the model increases monotonically as a function of ΔI. This again can be taken as an indicator (provided that fMRI signal saturation effects are minimized) that attractor decision-making processes of the type described here are taking place in a brain region. The percentage change in Fig. 4b was calculated by convolution of the firing rates of the neurons in the D1 and D2 populations. Similar effects, though smaller in degree, were found when the percentage change of the BOLD signal was calculated from the synaptic activity in all populations of neurons in the network (D1, D2, GABA, and non-specific, see Fig. 1), as shown in Supplementary Material Fig. 1a. Interestingly, the percent change in the BOLD signal is approximately linearly related throughout this range to ΔI (for Fig. 4b $r = 0.994$, $p < 10^{-7}$; for Fig. SM 1a $r = 0.991$, $p < 10^{-6}$). The effects shown in Figs. 4a and b can be related to the earlier onset of a high firing rate attractor states when ΔI is larger (see Figs. 2 and 3b), and to a higher firing rate when in the attractor states (as shown in Figs. 2 and 3a). As expected from the decrease in the variability of the neuronal activity as ΔI increases (Fig. 3a), the variability (standard deviation) in the predicted BOLD signal also decreases as ΔI increases, as shown in Fig. 4b ($r = 0.955$, $p < 10^{-4}$).

fMRI signals linearly related to the easiness of decisions, and to decision confidence

Fig. 5a shows experimental data with the fMRI BOLD signal measured on easy and difficult trials of the olfactory affective decision task (left) and the thermal affective decision task (right). The upper records are for prefrontal cortex medial area 10 in a region identified by criteria described in Methods as being involved in choice decision making. (The criteria were that a brain region for identical stimuli should show more activity when a choice decision was being made than when a rating on a continuous scale of affective value was being made.) Fig. 5a shows for medial prefrontal cortex area 10 that there is a larger BOLD signal on easy than on difficult trials. The top diagram shows the medial prefrontal area activated in this contrast for decisions about which olfactory stimulus was more pleasant (yellow) and for decisions about whether the thermal would be chosen in future based on whether it was pleasant or unpleasant (red).

In more detail, for the thermal stimuli, the contrast was the warm2 and the cold trials (which were both easy in that the percentage of the choices were far from the chance value of 50%, and in particular were 96 ± 1% (mean ± sem) for the warm, and 18 ± 6% for the cold), versus the mixed stimulus of warm2 + cold (which was difficult in that the percent of choices of "Yes, it would be chosen in future" was 64 ± 9%). For the temperature easy vs. difficult decisions about pleasantness, the activation in medial area 10 had peaks at [4 42 −4] $z = 3.59$ $p = 0.020$ and [6 52 −4] $z = 3.09$ $p = 0.045$.

For the olfactory decision task, the activations in medial area 10 for easy vs. difficult choices were at [−4 62 −2] $z = 2.84$ $p = 0.046$, confirmed in a finite impulse response (FIR) analysis with a peak at 6-8 s after the decision time [−4 54 −6] $z = 3.50$ $p = 0.002$. In the olfactory task, the easy trials were those in which one of the pair of odors was from the pleasant set, and the other from the unpleasant set. The mean difference in pleasantness, corresponding to ΔI, was 1.76 ± 0.25 (mean ± sem). The difficult trials were those in which both odors on a trial were from the pleasant set, or from the unpleasant set. The mean difference in pleasantness, corresponding to ΔI, was 0.72 ± 0.16. For easy trials, the percentage correct was 90 ± 2, and for difficult trials was 59 ± 8). No other significant effects in the a priori regions of interest (Grabenhorst et al., 2008; Rolls et al., 2010) were found for the easy vs. difficult trial contrast in either the thermal or olfactory reward decision task.

The lower records in Fig. 5a are for the same easy and difficult trials, but in parts of the pregenual cingulate and mid-orbitofrontal cortex implicated by the same criteria in representing the subjective reward value of the stimuli, but not in making choice decisions between the stimuli. For the pregenual cingulate cortex, there was a correlation of the activations with the subjective ratings of pleasantness of the thermal stimuli at ([4 38 −2] $z = 4.24$ $p = 0.001$). For the mid-orbitofrontal cortex, there was a correlation of the activations

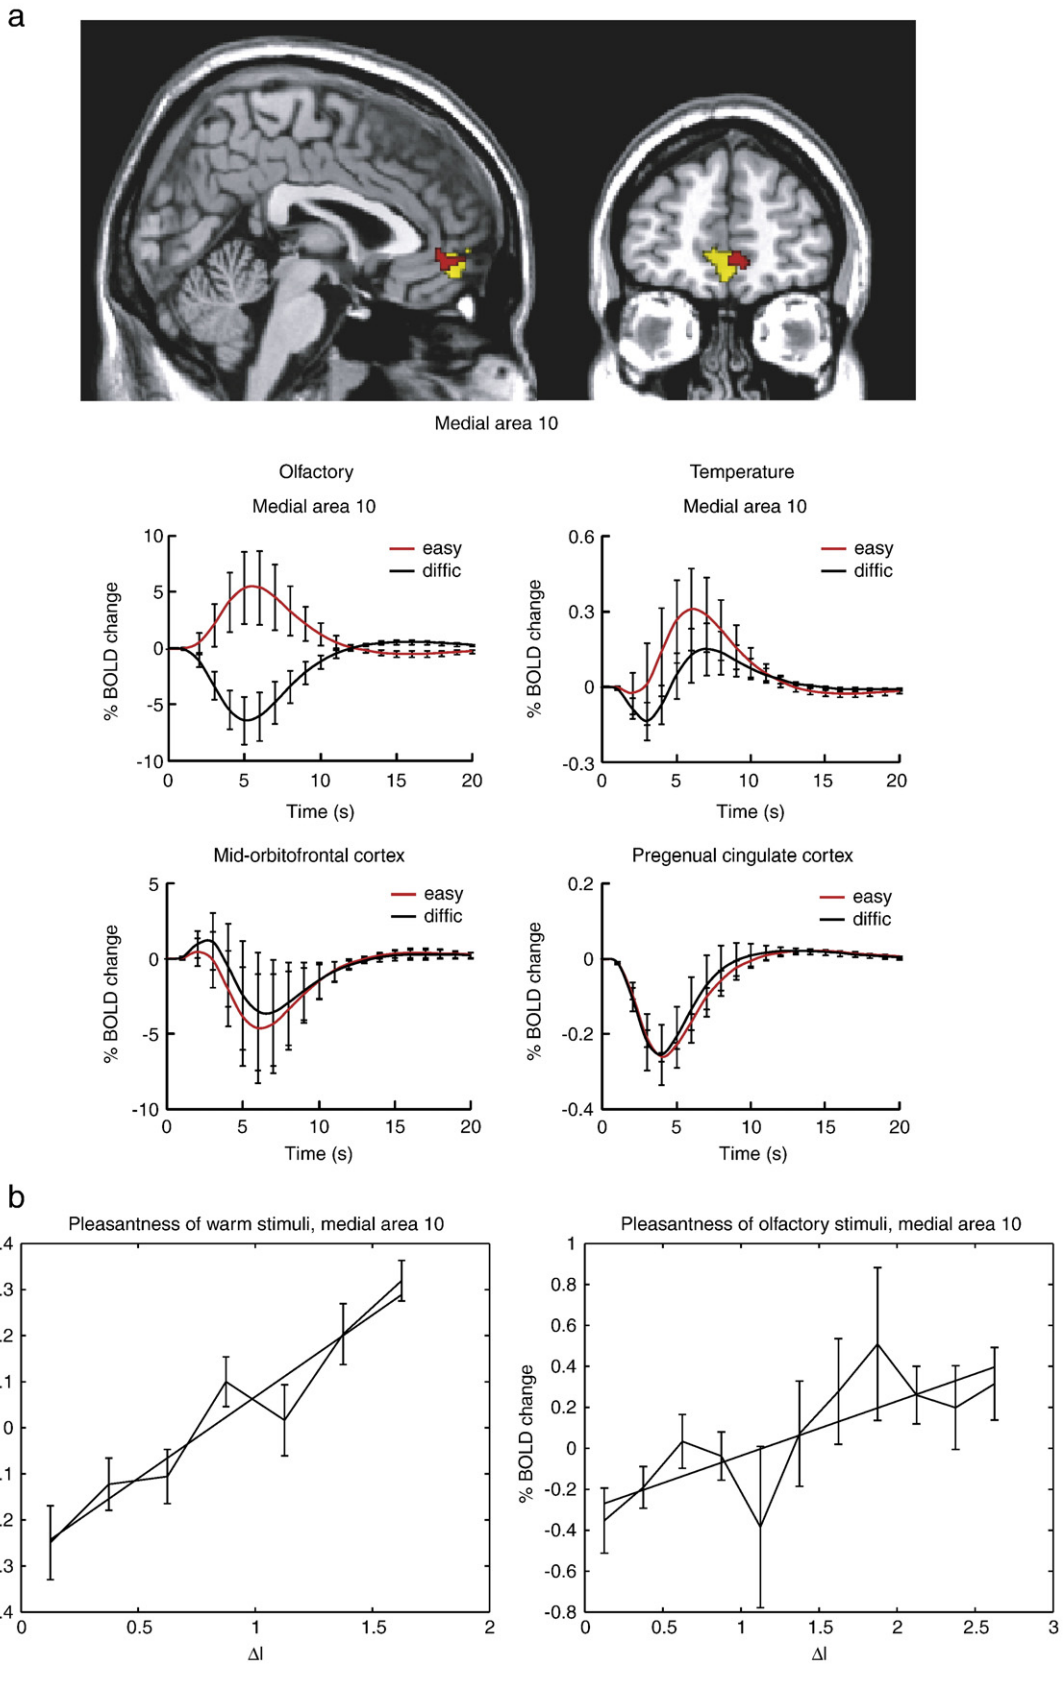

with the subjective ratings of pleasantness of the thermal stimuli at ([40 36 −12] $z = 3.13$ $p = 0.024$). The BOLD signal was similar in these brain regions for easy and difficult trials, as shown in Fig. 5a, and there was no effect in the contrast between easy and difficult trials.

Fig. 5b shows the experimental fMRI data with the change in the BOLD signal indicated as a function of ΔI, the difference in pleasantness of warm stimuli or olfactory stimuli about which a decision was being made, for medial prefrontal cortex area 10. For the olfactory decision task, ΔI was the difference in pleasantness (for a given subject) between the mean pleasantness of the first odor and the mean pleasantness of the second odor between which a decision was being taken, about which was more pleasant (Rolls et al., 2010). It is shown in Fig. 5b that there was a clear and approximately linear relation between the BOLD signal and ΔI for the olfactory pleasantness decision-making task ($r = 0.77$, $p = 0.005$). The coordinates for these data were as given for Fig. 5a. For the warm decision task, ΔI was the difference in mean pleasantness for a given subject from 0 for a given thermal stimulus about which a decision was being made about whether it should or should not be repeated in future (Grabenhorst et al., 2008). It is shown in Fig. 5b that there was a clear and approximately linear relation between the BOLD signal and ΔI for the thermal pleasantness decision-making task ($r = 0.96$, $p < 0.001$). The coordinates for these data were as given for Fig. 5a. We did not attempt to measure whether the trial to trial variability of the BOLD responses decreases with increasing ΔI as predicted (Fig. 4b) because of the inherent measurement variability from trial to trial of the BOLD response, but this would be useful in neuronal recording studies.

Discussion

The attractor neuronal network model of decision-making described here provides a firm foundation for understanding fMRI BOLD signals related to decision-making. The model provides two reasons for the fMRI signal being larger on easy than on difficult trials during decision-making. The first is that the neurons have somewhat higher firing rates on easy trials (e.g. those with a large ΔI) than on difficult trials, as shown in Fig. 3b. The reason for the faster response is that the difference in the inputs to the two attractors modifies the flow landscape so that the activity has a smaller barrier over which to jump towards one of the attractors (see Fig. 1c and Supplementary Material). The fMRI BOLD signal thus starts sooner, even if by only a time in the order of 0.5 s, as shown in Fig. 3b. The second reason is that the firing rates become higher in the winning attractor as ΔI increases, as shown in Fig. 3a, and this also contributes to a larger BOLD signal for easy vs. difficult decisions. The firing rate of the winning attractor is higher with a large ΔI because the external input adds to the internal recurrent collateral synaptic effect (see Fig. 1a) to lead to a higher firing rate. This effect is increased because the external input to the D2 neurons is low, and so it competes less with the D1 attractor, also leading to higher rates in the D1 attractor with high ΔI. Thus this approach from integrate-and-fire networks provides a clear foundation for understanding the shorter neuronal response times (Kim and Shadlen, 1999; Shadlen and Newsome, 2001) and the larger BOLD signals on easy than on difficult trials found in brain regions (Heekeren et al., 2004; Heekeren et al., 2008) that, as shown here, include the medial prefrontal cortex area 10 (Fig. 5).

The present findings also validate a larger BOLD signal on easy than on difficult trials as an fMRI signature of a region involved in decision-making, in that the areas that show this effect are areas implicated by other criteria in decision-making, including their larger activation on trials on which choice decisions must be made, rather than continuous ratings of affective value (Grabenhorst et al., 2008; Rolls et al., 2010) (Fig. 5). The orbitofrontal and pregenual cortex, where continuous affective value but not choice is represented (Grabenhorst et al., 2008; Rolls and Grabenhorst, 2008; Rolls et al., 2010), do not show this easy vs. difficult effect in the BOLD signal, as shown by the experimental fMRI results in Fig. 5. This is what is predicted for brain areas not involved in binary choice decision-making and its emergent property confidence, but in representing pleasantness on a continuous scale, which is likely to be a precursor and prerequisite for choice decisions about pleasantness (Grabenhorst et al., 2008; Rolls and Grabenhorst, 2008; Rolls et al., 2010). Consistently, the orbitofrontal and pregenual cingulate cortex project to the medial prefrontal cortex area 10 (Price, 2006). We note that ΔI is not related in any simple way to the pleasantness of the second odor, presented at the onset time for the fMRI analyses described here. Instead, ΔI is the absolute value of the difference in the pleasantness of the first and second odors, and this could be high if the second odor is much less pleasant than the first. These are useful control findings.

The model and its analysis described here also make predictions about the type of neuronal activity that will be found in brain areas involved in the choice part of decision-making, in which a categorical state must be reached. The predictions are that the neuronal activity will dynamically progress to a state in which there is a high firing of one population of neurons that occurs when one choice is made, with the other populations showing low activity, often below the spontaneous level, due to inhibition from the other attractor (especially when ΔI is high, as shown in Fig. 2). This high vs. low neuronal activity with almost a binary distribution of firing rates (as the decision is reached, either firing with a high firing rate, or a low firing rate, that predicts which decision is reached) is an important characteristic of a brain region performing the type of decision-making described here. However, although the firing rates are categorical in this sense, there is predicted to be a small effect of ΔI on the firing rate when in the attractor, as shown in Fig. 3, with the additional prediction that the firing rate of the losing populations will tend to become smaller as ΔI increases, as shown in Fig. 3a. Another diagnostic property of the neuronal activity in this type of decision-making system is that the reaction times of the neurons (the time it takes to reach an attractor) will decrease as a function of ΔI, as shown in Fig. 3b. Another fundamental property of this type of decision-making network is that it inherently shows statistical variation, with neuronal noise from the almost random, Poisson, spike times for a given mean rate influencing the time when the system starts to fall into an attractor, and then a rapid change of firing rate, due to the supralinear positive feedback from the recurrent collateral synaptic connections (the recurrent collateral effect) starting to take off (Deco et al., 2009; Rolls and Deco, 2010). Some examples of the noisy temporal evolution of the firing rate on individual trials are shown in Fig. 2.

Given that it is well established that subjective decision confidence increases with discriminability, ΔI (Jonsson et al., 2005; Vickers, 1979; Vickers and Packer, 1982), the degree of confidence in a decision emerges from the model as being reflected by the firing rates of the

Fig. 5. (a) Top: Medial prefrontal cortex area activated on easy vs. difficult trials in the olfactory pleasantness decision task (yellow) and the thermal pleasantness decision task (red). Middle: experimental data showing the BOLD signal in medial area 10 on easy and difficult trials of the olfactory affective decision task (left) and the thermal affective decision task (right). This medial area 10 was a region identified by other criteria (see text) as being involved in choice decision making. Bottom: experimental data showing the BOLD signal for the same easy and difficult trials, but in parts of the pregenual cingulate and mid-orbitofrontal cortex implicated by other criteria (see text) in representing the subjective reward value of the stimuli on a continuous scale, but not in making choice decisions between the stimuli, or about whether to choose the stimulus in future. (b) Experimental fMRI data showing the change in the BOLD signal (mean ± sem, with the fitted linear regression line shown) as a function of ΔI, the difference in pleasantness of warm stimuli or olfactory stimuli about which a decision was being made, for medial prefrontal cortex area 10.

neurons involved in the choice decision-making, as shown in Fig. 3a (with the percent correct, which is closely correlated with subjective confidence (Jonsson et al., 2005; Vickers, 1979; Vickers and Packer, 1982), shown in Fig. 3c). (The higher firing rates when in the attractor on easy vs difficult trials are also clearly illustrated in Fig. 2). In essence, one's confidence in a decision increases as the decision becomes easier and the percentage correct becomes better (Jonsson et al., 2005; Vickers, 1979; Vickers and Packer, 1982), and this is indexed quantitatively here by increasing values of ΔI or discriminability, which produces higher firing rates of the winning attractor. In fact, as shown in Fig. 3a, the difference of firing rates between the winning and losing populations also clearly reflects the confidence in the decision. Thus decision confidence emerges as a property of this type of attractor network model of decision-making. Confidence does not emerge naturally from the Bayesian approach (Beck et al., 2008; Ma et al., 2006; Ma et al., 2008), which in any case is at the phenomenological level of description, rather than describing the neural mechanism. However, the approach described here is not inconsistent with the Bayesian approach, or with the accumulator or race models (again phenomenological, and not making predictions about synaptic, neuronal, and BOLD activity as described in the Supplementary Material) which describe a diffusion process (Ratcliff et al., 1999; Smith and Ratcliff, 2004; Vickers and Packer, 1982), for the attractor model in fact operates to implement a non-linear diffusion process during the decision-making (Roxin and Ledberg, 2008). The integrate-and-fire attractor model of decision-making we describe thus has the advantages that not only is it consistent with earlier diffusion models that provide a good account of behavioral data on decision-making, but it also describes a neural mechanism for the decision-making, makes neurobiological predictions some of which have been confirmed in this paper, and has a representation of confidence as an emergent property. We suggest that the extension of attractor models to multiple decisions (e.g. Furman and Wang, 2008; Albantakis and Deco, 2009) may also reflect confidence in a similar way. In work by Kepecs et al. (2008) in the rat orbitofrontal cortex it has been shown that there are negative outcome neurons (somewhat similar to the primate orbitofrontal cortex error neurons of Thorpe et al. (1983), and positive outcome neurons, both of which reflect ΔI. We account for these rat neurons elsewhere (Insabato et al., 2010; Rolls and Deco, 2010)) by a further network that receives continuous-valued confidence-related information from the type of network described here, and then takes a decision based on the continuous confidence value.

It is suggested that decision-making categorisation effects found in other brain areas [area LIP about global visual motion (Churchland et al., 2008; Gold and Shadlen, 2002, 2007; Shadlen and Newsome, 1996); dorsolateral prefrontal and ventral premotor cortex about the intensity of odors (Rolls et al., 2010); the ventral premotor cortex about the frequency of vibrotactile stimuli (Deco et al., 2009; Romo and Salinas, 2003)] operate by a similar noise-influenced attractor mechanism which effectively encodes decision confidence (Rolls and Deco, 2010).

Another interesting property of the decision-making model described and analyzed here, with quite clear neurophysiological and behavioural predictions, is that the percentage correct performance asymptotes to perfect at 100% with quite moderate levels of ΔI (Fig. 3c), whereas the neuronal responses continue to become higher as ΔI increases further (Fig. 3a), the BOLD signals become larger (Fig. 4b), and the reaction times decrease further (Fig. 3b). Further, the inherent trial-by-trial variability in human choice reaction times, the neuronal mechanism for which has always been an enigma, can now be understood in terms of statistical fluctuations in decision-making networks in the brain (Marti et al., 2008; Rolls and Deco, 2010).

There are many novel findings described in this paper. First, we show how the firing rates of the winning and losing attractors reflect the easiness of the decision, measured by ΔI. Although the neurons make what is essentially a binary choice, with one or the other population winning, we show here that the firing rates are not just high or low, but in fact show some increase with ΔI. This was not shown in our earlier work (Deco and Rolls, 2006). The second point follows from this, that confidence in a decision, which is well established to be greater for easy than for difficult decisions (Jonsson et al., 2005; Vickers, 1979; Vickers and Packer, 1982), is an emergent property of an integrate-and-fire attractor network model of decision making. The easiness of the decision, and thereby decision confidence, is reflected in the firing rates of the winning and the losing attractors. We note that this point about decision confidence has not been shown before (Deco and Rolls, 2006; Wang, 2002), and moreover that different models of decision-making, such as the accumulator (or race) models (Carpenter and Williams, 1995; Ratcliff et al., 1999; Smith and Ratcliff, 2004; Vickers, 1979; Vickers and Packer, 1982) make no predictions about how confidence emerges from the decision-making process and how this is related to neuronal firing rates, because they are not models of how decision-making is made by brain mechanisms, but are instead mathematical models of artificial processes. Third, we show that the firing rates of the neurons of the winning attractor (and the total synaptic currents) increase approximately linearly with ΔI, and of the neurons of the losing attractor decrease approximately linearly with ΔI. This is a new finding. A fourth discovery is that we then use the integrate-and-fire attractor model to predict the fMRI BOLD signal, and show that the net effect of increasing ΔI is to increase the BOLD signal, and indeed to do so linearly with increasing ΔI. The effects in the winning attractor thus dominate the predicted BOLD signal. This is a new finding. A fifth new finding is that we show that the same prediction for the BOLD signal as a function of ΔI is made from the firing rates of the neurons and from the synaptic currents. This itself is interesting, for it shows that in a model of neocortical effects, the neuronal firing rates and the total currents in the neurons are closely (linearly) related to each other. The reason for this is that the neocortex operates primarily as an excitatory system, with excitatory inputs to a cortical area, and excitatory effects between neurons in an area, so that when overall the induced synaptic currents increase, so do the neuronal firing rates (Rolls, 2008). The inhibition is largely local, and serves to control the excitation. The sixth discovery is that when decisions about pleasantness are made, the BOLD signal increases approximately linearly with ΔI in a brain area implicated by other findings in this type of choice, the medial prefrontal cortex. This relation with ΔI has not been investigated before, as far as we know. This discovery was made in two separate fMRI experiments, one involving choices between the pleasantness of thermal stimuli applied to the hand, and the other involving choices between the pleasantness of olfactory stimuli. A seventh novel contribution was that this increasing BOLD signal with ΔI was not found in medial orbitofrontal cortex areas not implicated in making choices, but instead implicated by correlations with subjective pleasantness ratings in representing the continuous-valued representations of affective value that are the inputs to a choice decision process between two stimuli with different affective values. This is an important finding because it helps to define the distinct functional contributions of different areas in the prefrontal cortex to value-based decision-making. We note that this is one of the first reports of a functional dissociation between different areas within the ventromedial prefrontal cortex. This brain region is strongly implicated in decision-making by fMRI studies but very few studies have been able to dissect the functional roles of its anterior and posterior sub-regions in decision-making. Findings six and seven provide powerful validation of this theory of decision-making. In future investigations, it will be of interest to investigate with the model and theoretically how the model performs on error trials, what predictions the model makes about the fMRI signal on error trials, and whether the predicted effects are found in fMRI investigations (Rolls, Grabenhorst and Deco, in preparation).

We note that in earlier analyses of the neuroimaging data of the two studies utilized here (Grabenhorst et al., 2008; Rolls et al., 2010) we focused on a contrast of trials where decisions about rewards were made vs. trials where the rewards were evaluated in the absence of decision-making. The hypotheses for these analyses and the contrasts tested in the fMRI analyses were not informed by computational modeling results. These analyses led to the identification of the medial prefrontal cortex area 10 as a main brain area involved in decision-making. In the present paper we go beyond this analysis approach and address different questions. We used a biologically plausible computational model of decision-making to derive quantitative hypotheses about the types of neural signatures that should be found in a decision-making brain area. We then used the decision trials from the above studies to test these hypotheses. We did this by first performing new analyses on the behavioral data of the studies to derive subject-specific contrasts and regressors (including easy vs. difficult choices, and ΔI), and then used these as regressors for new contrast and regression analyses of the fMRI data. A novel contribution of the present paper with respect to our earlier reports is that we define a neural signature for a choice system in the brain and then show that this signature is present in the medial prefrontal cortex area identified in our previous studies. By combining the data from our two previous studies we are able to show that this neural signature is present both when different types of reward-based decisions are made and when decisions about different types of reward are made, thus providing strong support for our conclusions.

Overall, a major strength of the current approach is that the model incorporates effects from many levels, including the synaptic and biophysical levels, and the abstract level of statistical fluctuations in dynamical systems, and is able to make predictions all the way from synaptic currents, to neuronal activity, to fMRI signals, and to behavioral choice and subjective confidence, and offers a quantitative approach to our understanding at all these levels (Rolls, 2008; Rolls and Deco, 2010). It is in this sense a mechanistic approach, which takes into account details of the underlying biophysical and network mechanisms in the brain, and then can account for behavioral data, even to the level of accounting for the confidence in a decision, which arises as an emergent property of the mechanism (Rolls and Deco, 2010). Moreover, the approach generalizes to other processes implemented by attractor networks in the brain including memory recall, and indeed we propose that confidence in a recalled memory when the recall is correct or in error is an emergent property of the memory recall process (Rolls, 2008; Rolls and Deco, 2010) in the same way as that described here.

Acknowledgments

F.G. was supported by the Gottlieb-Daimler- and Karl Benz-Foundation and by the Oxford Centre for Computational Neuroscience. G.D. received support from the McDonnell Centre for Cognitive Neuroscience at Oxford University. The fMRI investigation was performed at the Centre for Functional Magnetic Resonance Imaging of the Brain (FMRIB) at Oxford University, and we thank Peter Hobden, Siri Leknes, Katie Warnaby, and Irene Tracey for their help.

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at doi:10.1016/j.neuroimage.2010.06.073.

References

Abeles, M., 1991. Corticonics—neural circuits of the cerebral cortex. Cambridge University Press, New York.
Albantakis, L., Deco, G., 2009. The encoding of alternatives in multiple-choice decision making. Proc. Natl Acad. Sci. USA 106, 10308–10313.
Beck, J.M., Ma, W.J., Kiani, R., Hanks, T., Churchland, A.K., Roitman, J., Shadlen, M.N., Latham, P.E., Pouget, A., 2008. Probabilistic population codes for Bayesian decision making. Neuron 60, 1142–1152.
Behrens, T.E., Woolrich, M.W., Walton, M.E., Rushworth, M.F., 2007. Learning the value of information in an uncertain world. Nat. Neurosci. 10, 1214–1221.
Brunel, N., Wang, X.J., 2001. Effects of neuromodulation in a cortical network model of object working memory dominated by recurrent inhibition. J. Comput. Neurosci. 11, 63–85.
Carpenter, R.H.S., Williams, M.L., 1995. Neural computation of log likelihood in control of saccadic eye movements. Nature 377, 59–62.
Churchland, A.K., Kiani, R., Shadlen, M.N., 2008. Decision-making with multiple alternatives. Nat. Neurosci. 11, 693–702.
Collins, D.L., Neelin, P., Peters, T.M., Evans, A.C., 1994. Automatic 3D intersubject registration of MR volumetric data in standardized Talairach space. J. Comput. Assist. Tomogr. 18, 192–205.
Deco, G., Rolls, E.T., 2005. Attention, short-term memory, and action selection: a unifying theory. Prog. Neurobiol. 76, 236–256.
Deco, G., Rolls, E.T., 2006. Decision-making and Weber's Law: a neurophysiological model. Eur. J. Neurosci. 24, 901–916.
Deco, G., Scarano, L., Soto-Faraco, S., 2007. Weber's law in decision making: integrating behavioral data in humans with a neurophysiological model. J. Neurosci. 27, 11192–11200.
Deco, G., Rolls, E.T., Romo, R., 2009. Stochastic dynamics as a principle of brain function. Prog. Neurobiol. 88, 1–16.
Friston, K.J., Worsley, K.J., Frackowiak, R.S.J., Mazziotta, J.C., Evans, A.C., 1994. Assessing the significance of focal activations using their spatial extent. Hum. Brain Mapp. 1, 214–220.
Friston, K.J., Glaser, D.E., Henson, R.N., Kiebel, S., Phillips, C., Ashburner, J., 2002. Classical and Bayesian inference in neuroimaging: applications. Neuroimage 16, 484–512.
Furman, M., Wang, X.J., 2008. Similarity effect and optimal control of multiple-choice decision making. Neuron 60, 1153–1168.
Genovese, C.R., Lazar, N.A., Nichols, T., 2002. Thresholding of statistical maps in functional neuroimaging using the false discovery rate. Neuroimage 15, 870–878.
Glimcher, P.W., 2003. The neurobiology of visual-saccadic decision making. Annu. Rev. Neurosci. 26, 133–179.
Gold, J.I., Shadlen, M.N., 2002. Banburismus and the brain: decoding the relationship between sensory stimuli, decisions, and reward. Neuron 36, 299–308.
Gold, J.I., Shadlen, M.N., 2007. The neural basis of decision making. Annu. Rev. Neurosci. 30, 535–574.
Goldman-Rakic, P.S., 1995. Cellular basis of working memory. Neuron 14, 477–485.
Grabenhorst, F., Rolls, E.T., Parris, B.A., 2008. From affective value to decision-making in the prefrontal cortex. Eur. J. Neurosci. 28, 1930–1939.
Hampton, A.N., O'Doherty, J.P., 2007. Decoding the neural substrates of reward-related decision making with functional MRI. Proc. Natl Acad. Sci. USA 104, 1377–1382.
Heekeren, H.R., Marrett, S., Bandettini, P.A., Ungerleider, L.G., 2004. A general mechanism for perceptual decision-making in the human brain. Nature 431, 859–862.
Heekeren, H.R., Marrett, S., Ungerleider, L.G., 2008. The neural systems that mediate human perceptual decision making. Nat. Rev. Neurosci. 9, 467–479.
Insabato, A., Pannunzi, M., Rolls, E.T., Deco, G., 2010. Confidence-related decision-making. J. Neurophysiol. 104, 539–547.
Jonsson, F.U., Olsson, H., Olsson, M.J., 2005. Odor emotionality affects the confidence in odor naming. Chem. Senses 30, 29–35.
Kepecs, A., Uchida, N., Zariwala, H.A., Mainen, Z.F., 2008. Neural correlates, computation and behavioural impact of decision confidence. Nature 455, 227–231.
Kiani, R., Shadlen, M.N., 2009. Representation of confidence associated with a decision by neurons in the parietal cortex. Science 324, 759–764.
Kiebel, S.J., Poline, J.B., Friston, K.J., Holmes, A.P., Worsley, K.J., 1999. Robust smoothness estimation in statistical parametric maps using standardized residuals from the general linear model. NeuroImage 10, 756–766.
Kim, J.N., Shadlen, M.N., 1999. Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. Nat. Neurosci. 2, 176–185.
Knutson, B., Rick, S., Wimmer, G.E., Prelec, D., Loewenstein, G., 2007. Neural predictors of purchases. Neuron 53, 147–156.
Ma, W.J., Beck, J.M., Latham, P.E., Pouget, A., 2006. Bayesian inference with probabilistic population codes. Nat. Neurosci. 9, 1432–1438.
Ma, W.J., Beck, J.M., Pouget, A., 2008. Spiking networks for Bayesian inference and choice. Curr. Opin. Neurobiol. 18, 217–222.
Marsh, A.A., Blair, K.S., Vythilingam, M., Busis, S., Blair, R.J., 2007. Response options and expectations of reward in decision-making: the differential roles of dorsal and rostral anterior cingulate cortex. Neuroimage 35, 979–988.
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.
Price, J.L., 2006. Connections of orbital cortex. In: Zald, D.H., Rauch, S.L. (Eds.), The Orbitofrontal Cortex. Oxford University Press, Oxford, pp. 39–55.
Ratcliff, R., Van Zandt, T., McKoon, G., 1999. Connectionist and diffusion models of reaction time. Psychological Reviews 106, 261–300.
Rolls, E.T., 2008. Memory, Attention, and Decision-Making: A Unifying Computational Neuroscience Approach. Oxford University Press, Oxford.
Rolls, E.T., Deco, G., 2002. Computational Neuroscience of Vision. Oxford University Press, Oxford.
Rolls, E.T., Deco, G., 2010. The Noisy Brain: Stochastic Dynamics as a Principle of Brain Function. Oxford University Press, Oxford.
Rolls, E.T., Grabenhorst, F., 2008. The orbitofrontal cortex and beyond: from affect to decision-making. Prog. Neurobiol. 86, 216–244.
Rolls, E.T., Grabenhorst, F., Parris, B.A., 2010. Neural systems underlying decisions about affective odors. J. Cogn. Neurosci. 22, 1069–1082.
Romo, R., Salinas, E., 2003. Flutter discrimination: neural codes, perception, memory and decision making. Nat. Rev. Neurosci. 4, 203–218.
Romo, R., Hernandez, A., Zainos, A., 2004. Neuronal correlates of a perceptual decision in ventral premotor cortex. Neuron 41, 165–173.

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.
Shadlen, M.N., Newsome, W.T., 1996. Motion perception: seeing and deciding. Proc. Natl Acad. Sci. USA 93, 628–633.
Shadlen, M.N., Newsome, W.T., 2001. Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. J. Neurophysiol. 86, 1916–1936.
Smith, P.L., Ratcliff, R., 2004. Psychology and neurobiology of simple decisions. Trends Neurosci. 27, 161–168.
Stephan, K.E., Weiskopf, N., Drysdale, P.M., Robinson, P.A., Friston, K.J., 2007. Comparing hemodynamic models with DCM. Neuroimage 38, 387–401.
Sugrue, L.P., Corrado, G.S., Newsome, W.T., 2005. Choosing the greater of two goods: neural currencies for valuation and decision making. Nat. Rev. Neurosci. 6, 363–375.
Thorpe, S.J., Rolls, E.T., Maddison, S., 1983. Neuronal activity in the orbitofrontal cortex of the behaving monkey. Exp. Brain Res. 49, 93–115.
Usher, M., McClelland, J.L., 2001. The time course of perceptual choice: the leaky, competing accumulator model. Psychological Reviews 108, 550–592.
Vickers, D., 1979. Decision Processes in Visual Perception. Academic Press, New York.
Vickers, D., Packer, J., 1982. Effects of alternating set for speed or accuracy on response time, accuracy and confidence in a unidimensional discrimination task. Acta Psychol. 50, 179–197.
Wang, X.J., 2002. Probabilistic decision making by slow reverberation in cortical circuits. Neuron 36, 955–968.
Wang, X.J., 2008. Decision making in recurrent neuronal circuits. Neuron 60, 215–234.
Worsley, K.J., Marrett, P., Neelin, A.C., Friston, K.J., Evans, A.C., 1996. A unified statistical approach for determining significant signals in images of cerebral activation. Hum. Brain Mapp. 4, 58–73.