Mammalian choices: combining fast-but-inaccurate and slow-but-accurate decision-making systems

Pete C. Trimmer^{1,*}, Alasdair I. Houston², James A. R. Marshall¹,
Rafal Bogacz¹, Elizabeth S. Paul³, Mike T. Mendl³ and John M. McNamara⁴

¹*Department of Computer Science, University of Bristol, Woodland Road, Bristol BS8 1UB, UK*

²*School of Biological Sciences, University of Bristol, Woodland Road, Bristol BS8 1UG, UK*

³*Department of Clinical Veterinary Science, University of Bristol, Langford House, Langford BS40 5DU, UK*

⁴*School of Mathematics, University of Bristol, University Walk, Bristol BS8 1TW, UK*

Empirical findings suggest that the mammalian brain has two decision-making systems that act at different speeds. We represent the faster system using standard signal detection theory. We represent the slower (but more accurate) cortical system as the integration of sensory evidence over time until a certain level of confidence is reached. We then consider how two such systems should be combined optimally for a range of information linkage mechanisms. We conclude with some performance predictions that will hold if our representation is realistic.

Keywords: neural pathways; decision making; signal detection theory; sequential probability ratio test; optimal foraging; mammalian brain

1. INTRODUCTION

We consider the neural pathways of mammalian decision making in the context of an animal attempting to forage while avoiding predators. There is considerable evidence that the danger posed by predators has a strong effect on the behaviour of their prey (see Lima (1998) for a review). Much of the theoretical work on this topic has focused on understanding how various factors influence the optimal behaviour of prey (e.g. Houston & McNamara 1999; Brown & Kotler 2004). Our aim is different, in that we are concerned with the optimal use of the neural mechanisms involved in responding to the threat of predation. Using parameters relating to brain structure and function, we construct and analyse a decision-making system based on two subsystems. One subsystem is fast while the other is more accurate, resulting in a trade-off between the use of the two components.

A novel aspect of our approach is that we find the optimal way to use these subsystems given various assumptions about the interactions between them. This means that we are able to predict how the presence of one neural system can shape the evolution of a new neural system.

2. BASIS FOR MODELLING ASSUMPTIONS

Experimental evidence suggests that mammalian species have evolved at least two distinct neural pathways for detecting and responding to signals of threat (LeDoux 1996; Morris *et al.* 1999; Zald 2003; Ohman 2005). Fear conditioning and lesioning studies of rats, and neuroimaging studies of humans, have indicated that the

amygdala plays a central role in the processing of threat-related sensory information, and the activation of defensive responses. For example, Vuilleumier *et al.* (2003) considered the processing of fear-related visual stimuli (faces) in humans and showed that shorter subcortical pathways appear to provide the amygdala with coarse but rapid sensory information, while longer cortical pathways provide more detailed information over a longer time-scale.

Fear responses can be produced from visual triggers, which are sufficiently brief that the subject is not consciously aware of the image shown (e.g. Ohman & Soares 1994), indicating that the triggers may be processed and acted upon purely subcortically. On the other hand, it has been demonstrated that cortical regions are capable of integrating sensory evidence over time to increase accuracy of judgments (Schall 2001; Shadlen & Newsome 2001; Cisek & Kalaska 2005; Yang & Shadlen 2007). In particular, in a task in which a monkey has to make a decision about a direction of motion from a noisy visual stimulus, the neurons in the frontal and parietal areas that correspond to the correct alternative gradually increase their firing rate (Shadlen & Newsome 2001). This suggests that these neurons accumulate input from sensory neurons, which are selective for the corresponding alternative. Furthermore, it has been reported that when the level of activity of the integrating neurons reaches a particular threshold, the choice is made and the corresponding action is initiated (Roitman & Shadlen 2002).

The accumulation of sensory evidence is also evident in behavioural data: the distributions of reaction times from choice tasks between two alternatives are very well described by a diffusion model assuming that the difference between evidence for the two alternatives is integrated until it reaches a certain threshold (Laming

* Author for correspondence (trimmer@compsci.bristol.ac.uk).

Electronic supplementary material is available at <http://dx.doi.org/10.1098/rspb.2008.0417> or via <http://journals.royalsociety.org>.

1968; Ratcliff 1978; Ratcliff & Smith 2004). The diffusion model is thus consistent with the above neurophysiological data (Smith & Ratcliff 2004; Bogacz 2007). Furthermore, it has been shown that neural network models can implement the diffusion model for certain parameter values (Usher & McClelland 2001; Bogacz *et al.* 2006; Wong & Wang 2006).

The central nucleus of the amygdala acts as an interface between sensory inputs and response control systems, triggering physiological and behavioural fear reactions. LeDoux (1996), who identified subcortical and cortical routes associated with auditory fear conditioning, suggested that the evolutionarily older, subcortical route has persisted despite the newer, more accurate cortical route, because it maintains an advantage in allowing individuals to respond more rapidly in life-threatening situations. Although these pathways have now been studied in some detail, the actual processes by which decisions (e.g. to flee) are made in response to stimuli remain unclear.

By measuring response times with different background probabilities of outcome, Carpenter & Williams (1995) have produced empirical evidence for two decision-making systems that operate at different speeds, but do not suggest how these are implemented.

Such findings motivate us to assume that the mammalian brain has two decision-making components and that these components are able to process data in different ways in a variety of contexts, including the detection of predators. One system is assumed to process a single stimulus very fast, to reach a binary decision without waiting for additional information; we refer to this as the 'thalamic system'. The second system is assumed to be capable of reaching the same binary decisions that the first system can reach but has, at each stage, an additional option of waiting for more information; we refer to this as the 'cortical system'.

3. GENERAL SCENARIO AND APPROACH

We analyse a hypothetical situation in which an animal forages continuously unless it decides that there is sufficient risk of a predator being present that it should leave to find a safe area. The animal can choose to leave at any time. If the probability of a predator being present is perceived to be small, the focal animal may never choose to leave and would then forage indefinitely (unless it is killed).

The optimal decision at any stage depends upon the value of continuous successful foraging, v_f , the cost of taking anti-predator action, c and the probability of survival in the presence of a predator if anti-predator action is taken, $S(T)$, where the random variable T is the time to decide to take anti-predator action. We assume that if the animal leaves the area, then by paying the cost c , it is able to reach an area in which it can continue to forage without risk of predation. The resulting pay-offs, in terms of reproductive value, are shown in table 1.

The scenario is assumed to start with a single stimulus (such as a sense of movement from peripheral vision, or the sound of a twig snapping), to which the thalamic system is capable of making an immediate response. The cortical system is also assumed to start its assessment at

Table 1. General pay-off matrix.

	forage	anti-predator
no predator	v_f	$v_f - c$
predator	0	$S(T)(v_f - c)$

Figure 1. Information gain: (a) shows the discrete signal distributions relating to the thalamic system and (b) illustrates the random walk nature of information gain with time in the cortical system.

the same time, and may cause additional data to be gathered rapidly (e.g. through visual saccades) before reaching a decision.

We assume that the evolutionary history and learning of the animal allows it to correctly estimate the probability, p , of a predator being present prior to the initial stimulus, and that there is only one possible predator so the probability of there being no predator is $(1 - p)$.

We first analyse the thalamic and cortical systems independently, before considering a number of ways in which the systems could be combined.

4. THALAMIC DECISION: SIGNAL DETECTION THEORY

We assume that the signal, x , received by the thalamic system via sensory inputs, is normally distributed according to $N(\mu_{\text{thal}1}, \sigma_{\text{thal}})$ or $N(\mu_{\text{thal}0}, \sigma_{\text{thal}})$, depending upon whether a predator is present or not, respectively, as shown in figure 1a.

The optimal stand-alone performance of the thalamic system can be calculated using standard signal detection theory (Egan 1975). Using the pay-off matrix of table 1, with $s = S(0)$ being the probability of survival in the presence of a predator if immediate anti-predator action is taken, we find that the optimal critical threshold for signal level is given by

$$\frac{f_{\text{thal}1}(x_{\text{threshold}})}{f_{\text{thal}0}(x_{\text{threshold}})} = \frac{(1-p)}{p} \frac{(v_f - (v_f - c))}{s(v_f - c)} = \frac{(1-p)c}{ps(v_f - c)}.$$

The expected pay-off for the stand-alone case is then

$$\begin{aligned}
 & pE(\text{payoff}|\text{predator}) + (1-p)E(\text{payoff}|\text{no predator}) \\
 &= p \left(1 - D \left(\frac{x_{\text{threshold}} - \mu_{\text{thal1}}}{\sigma_{\text{thal}}} \right) \right) s(v_f - c) \\
 &+ (1-p) \left(D \left(\frac{x_{\text{threshold}} - \mu_{\text{thal0}}}{\sigma_{\text{thal}}} \right) v_f \right. \\
 &+ \left. \left(1 - D \left(\frac{x_{\text{threshold}} - \mu_{\text{thal0}}}{\sigma_{\text{thal}}} \right) \right) (v_f - c) \right) \\
 &= p \left(1 - D \left(\frac{x_{\text{threshold}} - \mu_{\text{thal1}}}{\sigma_{\text{thal}}} \right) \right) s(v_f - c) \\
 &+ (1-p) \left(v_f - c \left(1 - D \left(\frac{x_{\text{threshold}} - \mu_{\text{thal0}}}{\sigma_{\text{thal}}} \right) \right) \right),
 \end{aligned}$$

where $D(z)$ denotes the cumulative distribution function for the standard normal distribution, i.e. the probability that a sample from a standard normal distribution is less than z .

5. CORTICAL DECISION: SEQUENTIAL PROBABILITY RATIO TEST WITH ONE FINITE BOUNDARY

We assume that the focal animal continues to forage until the likelihood of a predator being present (given the evidence accumulated so far) is sufficiently high that it is best to leave. This scenario can be framed in terms of the sequential probability ratio test (SPRT), which can be implemented by the diffusion model (discussed above). The SPRT is a hypothesis test introduced by Wald (1945), which tests between two hypotheses, updating the relative likelihood of each as new data arrives until deciding (with some pre-defined error probability) in favour of one of the hypotheses. We use the test with only one finite boundary, so the animal will never reach the decision to forage permanently (i.e. regardless of new information). The test therefore results in one of two outcomes at each stage:

- (i) decide that the likelihood of a predator being present is sufficiently high that anti-predator action should be taken or
- (ii) decide to gather additional data (before, potentially, reaching outcome 1).

Assuming that the information received by the animal is independent from moment to moment, the SPRT is a statistically optimal test, in that it requires the minimum number of samples to reach a decision with particular error probabilities. Thus, the SPRT is the theoretical limit of animal performance if drift does not vary across trials. (Otherwise, a modification to the equations such as urgency signals is required, as shown by Ditterich (2006a,b).) Our assumption that the mammalian brain is able to implement the SPRT is also based upon evidence that it can be implemented in a simple cortical system (Gold & Shadlen 2002; Bogacz et al. 2006).

The amount of time required to assess the situation will be governed by the rate and quality of information and the probability of making errors. If the probability of false alarm given that there is no predator, α , is required to be close to zero, the amount of time required to assess the situation will be large. We assume that when there is a predator present, the probability of survival decreases

exponentially with the amount of time spent before taking anti-predator action. Thus, a trade-off is required between a fast inaccurate decision and a slow accurate one.

We apply the SPRT to the problem of evaluating predator presence, to look at the optimal speed-accuracy trade-off of the stand-alone performance of the cortex, before combining the cortical system with that of the thalamus. In this scenario, only one of the SPRT boundaries is finite, as the animal is assumed to pay no cost for continuing to accrue information when no predator is present, except through the possibility of accidentally deciding that a predator is present (by crossing the positive boundary after some time, despite negative drift).

We assume that the focal animal receives a sequence $x_1, x_2, \dots$ of signals. The presence or absence of a predator is assumed to remain constant over the time of the animal deciding whether or not to take anti-predator action. If no predator is present, each signal is assumed to be an independent, identically distributed random variable from a normal distribution with mean μ_{cort0} and variance σ^2 . If a predator is present, each signal is an independent, identically distributed random variable from a normal distribution with mean μ_{cort1} and variance σ^2 , with $\mu_{\text{cort0}} < \mu_{\text{cort1}}$.

After n observations, the information captured by the animal is summarized by $Z(n) = \sum_{i=1}^n q_i$, where $q_i = \ln(f_1(x_i)/f_0(x_i))$, and $f_0(x)$ denotes the probability density of signal x when no predator is present and $f_1(x)$ denotes the probability density of signal x when a predator is present.

As information is accumulated and n increases over time, the movement of $Z(n)$ is a biased random walk, as illustrated in figure 1b.

The mean drift and variance of the random walk in log space under each scenario (predator present or not) can be found directly from the two signal distributions per unit of information, as shown in the electronic supplementary material A. As the time between each datum is reduced to zero (and the corresponding information content with it so that the amount of information per unit time remains constant) and the information gain becomes a continuous process, the random walk becomes a Brownian motion (Wiener process) with constant drift. This is also shown in the electronic supplementary material A.

It is clear from figure 1b that the threshold governs both the expected decision time when a predator is present and the likelihood of false alarm when there is no predator. A small threshold value would result in fast decisions but a high false-alarm rate, and vice versa with a large threshold value. The optimal speed-accuracy trade-off can be obtained by first representing the probability of survival (when a predator is present) as a function of the false-alarm rate in the pay-off matrix, so $R = R(\alpha)$.

We assume that, *a priori*, the probability of a predator being present is p . Then,

$$\begin{aligned}
 E(\text{pay-off}) &= (1-p)((1-\alpha)v_f + \alpha(v_f - c)) + pR(\alpha)(v_f - c) \\
 &= (1-p)(v_f - \alpha c) + pR(\alpha)(v_f - c).
 \end{aligned}$$

To find the optimal value of α , we note that the expected pay-off is maximal when

$$\frac{dE(\text{pay-off})}{d\alpha} = -(1-p)c + p(v_f - c) \frac{dR}{d\alpha} = 0.$$

We first calculate the expected probability of survival, R , and its first derivative with respect to α .

Figure 2. Separation of cortical signal distribution means (per unit time) to match the thalamic performance. (a) With p being varied and $\mu_{\text{thal1}} - \mu_{\text{thal0}} = 2$. (b) With $p = 0.5$ while the thalamic distributions means are varied. In each case, $v_f = 10$, $s = 1 = c$ and all distributions have unit variance (per unit time for the cortical system).

We assume that the probability of survival, given that there is a predator, decreases with the decision time according to $e^{-\theta T}$, where T is a random variable describing the time to reach a decision, and that this time is a function of the position of the decision boundary. Note that the boundary position sets the probability of false alarm, α .

In electronic supplementary material B, we show that $R = s\alpha^{k_1}$ and the optimal value of α is given by

$$\alpha^{k_1-1} = \frac{(1-p)c}{p(v_f - c)k_1 s},$$

where s is the probability of survival if immediate action is taken in the presence of a predator and $k_1 = (-\mu + \sqrt{\mu^2 + 2\theta\eta^2})/\eta^2$.

This results in a final expected pay-off

$$E(\text{pay-off}) = (1-p) \left(v_f + c\alpha \left(\frac{1}{k_1} - 1 \right) \right).$$

6. INDEPENDENT CORTEX AND THALAMUS

We now consider the performance of the two systems when run independently. We assume that the cortical system is reset at the start of the scenario (upon the initial stimulation) and that the data gathered by the cortex are independent of the thalamic trigger, within a given context of a predator being present or not.

Figure 2a shows the critical separation of distribution means for the cortex to match the expected pay-off of the thalamic system as a function of p for a particular set of parameter values.

Varying $(\mu_{\text{thal1}} - \mu_{\text{thal0}})$ produces a family of curves, each following the trend of figure 2a. As the probability of a predator rises, the necessary separation in cortical distributions increases ever more sharply for the performance to match that of the thalamic system. This indicates that the thalamic system may be better able to deal with situations where the probability of a predator is great (the thalamic system achieves this by often deciding to run away immediately).

Figure 2b shows the separation of distribution means in the cortex required to match the expected pay-off for a range of thalamic distribution separations, when p is fixed and the other values are held according to figure 2a.

Figure 3. Schematic flow diagram of the combined system.

The near-exponential slope of the line of equivalent performance in figure 2b is caused by the modelling assumption that the decision for the thalamic system will be reached in zero time. Consequently, as the separation between the thalamic distributions increases and an ever smaller percentage of cases are mis-classified by the thalamic system, a significant increase in cortical separations is required to reach the same expected pay-off due to the cortical system paying a time penalty.

We now turn to treating the thalamic and cortical systems as two sub-components of a larger system.

7. COMBINING SYSTEMS

There are several ways in which the cortical and thalamic systems might be combined, with the go/no-go decision finally being governed by a central switch in some region of the brain, such as the amygdala. McHaffie *et al.* (2005) discuss why it makes sense for the architecture of the brain to incorporate a central switch and highlight the use of particular parts of the basal ganglia.

The general computational representation is shown in figure 3.

(a) Simple chaining

We first consider simply linking the two independent systems end-to-end. If the thalamic system does not make the decision to leave the area (so the animal continues to forage), the cortex starts to gather data and the animal will still decide to leave if the cortical process reaches that conclusion. We refer to this as the 'chained' system.

Figure 4. Optimized performance for different systems: (a) shows the expected pay-offs for each system (solid line, thalamus alone; triangles and dotted line, cortex alone; circles and dashed line, chained; crosses and solid line, thalamus aware; triangles and dashed line, cortex aware; dotted line, both aware) and (b) shows the optimal thalamic threshold values, measured with respect to the mean value of the distribution when no predator is present (dashed line, thalamus aware; dotted line, thalamus alone; solid line, both aware). In each case, $v_f=10$, $\mu_{\text{thal1}}-\mu_{\text{thal0}}=3$, cortical separation per unit time=8, $s=1=c$ and all distributions have unit variance (per unit time for the cortical system).

To calculate the expected pay-off, we first calculate the probability of the thalamus deciding to take anti-predator action:

$$P(\text{thal acts}) = pP(\text{anti-predator}|\text{predator}) \\ + (1-p)P(\text{anti-predator}|\text{no predator}).$$

The likelihood of a predator being present when the thalamic system has allowed foraging to continue is less than p because the likelihood is conditional upon the thalamic system having decided against anti-predator action. The resulting conditional probability of a predator being present when the cortex starts operating is

$$p' \equiv P(\text{predator}|\text{cortex used}) \\ = \frac{pD\left(\frac{x_{\text{threshold}} - \mu_{\text{thal1}}}{\sigma_{\text{thal}}}\right)}{pD\left(\frac{x_{\text{threshold}} - \mu_{\text{thal1}}}{\sigma_{\text{thal}}}\right) + (1-p)D\left(\frac{x_{\text{threshold}} - \mu_{\text{thal0}}}{\sigma_{\text{thal}}}\right)}. \quad (7.1)$$

The expected pay-off, given that the cortex is used, can now be calculated from the pay-off matrix using the actual probability of a predator, p' , but using the values of α and $R(\alpha)$, which were optimized using the original probability of a predator, p

$$E(\text{pay-off}|\text{cortex used}) = (1-p')(v_f - \alpha c) + p'R(\alpha)(v_f - c).$$

The overall expected pay-off for the system can therefore be calculated as

$$E(\text{pay-off}) = P(\text{thal acts})(v_f - c) \\ + (1 - P(\text{thal acts}))E(\text{pay-off}|\text{cortex used}).$$

Results for the simple combined system are shown as the chained line in figure 4a, which describes the performance in comparison with the stand-alone systems (and other systems yet to be discussed) for a particular set of parameters. For the figure, the chosen cortical separation provides a signal-to-noise ratio of 4, which corresponds to a fairly easy task for humans (see fig. 9b of Bogacz et al. submitted). The chosen separation of the thalamic distributions also corresponds to a relatively easy task for humans, towards the easiest of the tasks described by Maddox (2002).

Clearly, it is possible for the chained system to have a better performance than either subsystem acting independently, as can be seen from figure 4a when p is small. However, for a sufficiently high probability of a predator, the chained system performance is forced down to the pay-off of immediate anti-predator action. This is because if p is sufficiently large, the cortical system will set a zero threshold to take anti-predator action immediately, so the overall system will always take immediate anti-predator action no matter how separated the thalamic distributions. This can be seen by following the flow diagram of figure 3, which shows that if either subsystem were always to trigger immediate anti-predator action, then the overall system performance would correspond to that action.

(b) Cortical system 'aware' that the thalamic system has not acted

With simple chaining, the cortical system has been set up optimally for the case where the probability of a predator being present is p , but the thalamic system has already had the opportunity to act, so the actual (conditional) probability of a predator may alter significantly in the knowledge that the thalamic system has not acted.

If we assume that the cortex can use p' (calculated using equation (7.1)), rather than p , in the calculation of optimal α (and thus $R(\alpha)$) in the cortical system, a better overall expected pay-off is achieved than by the simple chained system. This is illustrated by the 'cortex-aware' line in figure 4a.

The overall performance of the cortex-aware system is better than the simple chained system or, for the parameter values of figure 4a, either subsystem operating independently. However, with poor thalamic separations (and thus performance), the overall system performance can be worse than the stand-alone cortex.

(c) Thalamic system aware that the cortical system can act

Knowing that the cortical system can still process incoming data if the thalamus does not invoke action, it can sometimes be beneficial for the thalamic threshold to be increased, to allow more of the ambiguous cases to be dealt with by the cortex.

Maximizing the expected pay-off by altering the thalamic threshold (with the cortex optimized as though it were stand alone), we obtain the ‘thalamus-aware’ line of figure 4a.

Figure 4b shows that the optimal thalamic threshold is higher when the thalamus is aware that the cortex could still make the decision to take anti-predator action. This can be viewed as the thalamus allowing the cortex to deal with some of the ambiguous cases near the (stand-alone) thalamic threshold, rather like the thalamus taking on the additional option of being able to wait for more information.

This system suffers from the same problem as the simple chained system in that if the cortex always takes immediate anti-predator action, the overall system performance is forced the pay-off relating to that action. For that reason, when the probability of a predator is sufficiently high, there is no optimal threshold value for the case of the thalamic system being aware of the cortical system, as shown in figure 4b.

This thalamus-aware-of-cortex case is arguably less relevant than the cortex-aware-of-thalamus case, as the thalamus is believed, from comparative studies, to be phylogenetically older than the cortex (MacLean 1982). However, understanding this thalamus-aware-of-cortex case is important to help understand the next system, where each subsystem is aware of the other, whether through evolutionary setting of parameters or direct information flow.

(d) *Each subsystem aware of the other*

Allowing each subsystem to be optimized with respect to the other results in a system, which outperforms the previous cases. The expected pay-off can be calculated using the equations for the cortex-aware case, with the additional maximization of expected pay-off across the possible threshold values. The results are shown in figure 4a,b as the ‘both-aware’ line.

The optimal thalamic threshold is still higher than the stand-alone case because the thalamus is aware that the cortex can still take action. However, the threshold is lower than the case of only the thalamus being aware of the cortex. This is because the cortex has a better knowledge of the probability of a predator (which is reduced from p to p'), so with the cortex being less prone to anti-predator action, the thalamic threshold is reduced to allow more actions to be taken by the thalamus.

Thus far, we have assumed that each subsystem may be aware of the other in terms of performance distributions and thresholds. However, it is also possible that the information used by the thalamus may also be available to the cortex. This is the final case that we consider.

(e) *Thalamus passes on likelihood estimate based upon signal level*

Rather than only knowing that the thalamic system has not fired, it is feasible that the cortex could have access to the signal datum (x) used by the thalamus. This datum could be used to give the cortex a more accurate estimate of the probability of a predator being present, which we shall call p''

$$p'' = \frac{pf_{\text{thal1}}(x)}{pf_{\text{thal1}}(x) + (1-p)f_{\text{thal0}}(x)}.$$

This can be used in the same cortical equations for α and R , just as p' was in the cortex-aware case. The alteration results in performance that is everywhere fractionally (but only

fractionally) better than the cortex-aware case. Consequently, there is no need to consider the both-aware system with this set-up, as we know that there will always be (only) a very small improvement.

We have not managed to prove that the performance discrepancy with the cortex-aware case will always be tiny, but we have a general line of reasoning to demonstrate that this will be the case. For this latest system (of the cortex receiving data relating to the thalamic trigger) to be a significant improvement upon the cortex-aware case, the cortex would have to gain significant additional information from the thalamic signal. The information gain would only be significant if the thalamic distributions were well separated—but if that were the case, then the thalamus would already be doing a good job of reaching the correct decision, so the effect of the cortex (and thus the change in performance) would be small. However, this argument would not apply if the variance of the two thalamic distributions were not equal; the data may then be considerably more useful to the cortex and thus the overall system.

Contrasting the cases of similar performance (of the cortex being aware that the thalamus has not fired, and the cortex receiving the thalamic information), there are two lines of reasoning, which differ over which might be more successful in a biological organism (given that they are each an improvement on the simple chained system). As the benefit of passing the thalamic information comes only with additional information flow and considerably more calculation of α and R (which has been assumed to occur in zero time for these results), it could be regarded as unlikely that such an adaptation would be beneficial in the real world. On the other hand, for the system to benefit from the cortex being optimized according to the thalamic distributions, the distributions themselves must be known (in some form) by the cortex. The answer therefore depends upon whether it is more advantageous to use additional memory or more computation.

8. EXPECTED DECISION TIME

The expected time for the cortex to decide upon anti-predator action, $E(T_c)$, can be calculated for the cases where that decision is reached as follows:

$$E(T_c|\text{AP action}) = \frac{P(\text{predator})E(T_c|\mu_+, \text{AP action}) + P(\text{no predator})E(T_c|\mu_-, \text{AP action})}{P(\text{predator})P(\text{AP action}|\mu_+) + P(\text{no predator})P(\text{AP action}|\mu_-)},$$

where ‘AP action’ denotes anti-predator action; and μ_+ and μ_- denote whether a predator is present or not, respectively, following the convention of the electronic supplementary material A. The denominator normalizes the estimate, because the animal will not always take anti-predator action if there is no predator present. With a predator present, only those decisions which are made before the animal is killed are taken into account (from a mathematical perspective, this amounts to the predator killing the focal animal instantaneously upon attack).

Electronic supplementary material C provides the necessary calculations for expected decision times and probabilities of anti-predator action. The expected decision times for the various system possibilities are shown in figure 5a.

Figure 5. Expected time to anti-predator action (if it occurs) against (a) probability of predator or (b) probability of action occurring. Squares and dashed line, cortex alone; solid line, thalamus aware; dashed line, chained; triangles and solid line, both aware; crosses and dotted line, cortex aware. In each case, $v_f = 10$, $\mu_{\text{thal1}} - \mu_{\text{thal0}} = 3$, cortical separation per unit time = 8, $s = c = \theta = 1$.

Contrasting the separation between cases in figure 4a with the relatively similar pay-offs of figure 4a, it is clear that, from an experimental perspective, it makes more sense to study reaction times than overall pay-offs when trying to distinguish which subsystem is prone to invoke action in a particular circumstance. Furthermore, the difficulty of setting and/or measuring the animal's perceived probability of a predator, p , can be largely bypassed by instead measuring the probability of anti-predator action and reaction times over a number of trials.

By varying the *a priori* probability of a predator and determining both the expected decision time when anti-predator action occurs (as in figure 5a) and the total probability of anti-predator action occurring, we obtain predictions that are essentially equivalent to those of figure 5a, but which rely upon axes which are more easily measured, as shown in figure 5b.

9. DISCUSSION

We have represented the decision-making process of the mammalian brain using two Bayesian subsystems, one dealing with discrete signals on an immediate basis and the other dealing with continuous data acquisition over time. While the parameter values we have used for plots in this paper are based upon realistic estimates from human studies (Maddox 2002; Bogacz *et al.* submitted), the trends we discuss apply across a wider range of values.

For very low probability that a predator is present, p , the expected pay-off tends towards the value of continuous foraging, v_f , whereas with high p , the expected pay-off is forced down to the reward for immediate predator avoidance, $v_f - c$. Therefore, the probability that a predator is present governs expected pay-off to a greater extent than the differences between subsystem linkage options, as indicated in figure 4a. As p increases, the faster (thalamic) subsystem will always tend to dominate decisions. This modelling prediction agrees with the empirical results of Carpenter & Williams (1995) who found that situations which occur with a higher probability are more likely to be dealt with by a high-speed processing system. The finding is also consistent with the work of Mobbs *et al.* (2007), who used functional magnetic resonance imaging to study the effect of threat level upon the activity of various regions of the human brain. The study found that as threat levels increased, 'brain activity

shifted from the ventromedial prefrontal cortex to the periaqueductal grey', i.e. the cortex has a more significant role when there is a low level of threat.

In behavioural ecology, studies typically focus upon outcomes rather than mechanisms (but see Real (1992, 1993) and Dukas (1998)). To truly understand the outcomes (including probably phylogeny, when comparative studies are not available), we need to understand the mechanisms, and these can sometimes be difficult to distinguish in terms of overall outcome (expected pay-off). For small p , there is little difference in overall pay-offs between the different linkage mechanisms, as indicated in figure 4a. However, situations with low probability can be numerous, so scaling the difference in performance by the number of occurrences an animal may face, the small differences can create significant fitness pressures (cf. Houston & McNamara 1986, 1999). From the perspective of empirical testing of the model, although the overall pay-off does not vary significantly across the linkage options for low p , the cases may be discriminable by measuring latency, as indicated in figure 5a.

For high p , the lines of figure 4a are fundamentally split into two groups. For those systems in which the cortex is not aware of the thalamus, the cortical system will always take immediate anti-predator action for a range of p values close to 1. This is because for any non-zero threshold, the variability of the Brownian motion process can have a significant effect upon the probability of survival. When the cortex is aware of the thalamic involvement, the cortex need not be so prone to taking action because the thalamic system will have dealt with many of the cases in which a predator is present. Consequently, if our model is representative of real mammalian brains, the thalamic subsystem will not atrophy because it provides a fitness benefit that the cortical system cannot subsume. This agrees with the suggestion of LeDoux (1996), which the evolutionarily older (thalamic) route has persisted because it maintains an advantage in allowing individuals to respond more rapidly in life-threatening situations.

When each subsystem is aware of the other, the optimal performance using these systems is attained; this will always result in a higher expected pay-off than either system operating alone.

The extent to which each subsystem may be aware of the other and able to optimize performance based upon expected performance of other subsystems is not yet clear.

However, Sole *et al.* (2003) provide evidence that the actions of pigeons (Columbidae) can be accurately described by a signal detection model that maximizes perceived reward; the birds are able to set more than one threshold, with each being based upon expected performance in relation to other thresholds. Similar studies of mammalian decisions (e.g. comparing humans and dolphins; *Tursiops truncatus*: Smith *et al.* 1995) have been carried out for certain kinds of task, again finding that expected performance levels guide decisions. These findings agree with the notion of the thalamic threshold being modified appropriately when another system is available as a back-up.

Although we have associated functions with particular parts of the brain in this paper, animals without a particular component may be able to reach similar outcomes using alternative parts and processes. For instance, Chittka *et al.* (2003) showed that (even without the collective intelligence of a group) bees are able to trade-off speed against accuracy.

Our results predict that actions or situations, which occur infrequently, will tend to be dealt with by the cortex rather than the thalamic system (as indicated in figure 5). It is then natural to ask whether the cortical component of decision making has evolved due to the fitness benefit of being able to deal with events which occur with relatively low probability, or to deal with common but subtle situations which require significant amounts of data, such as social interactions. The latter possibility is related to the alternative explanation that the cortex developed to monitor threat probabilities and to keep the thalamic system updated with estimates, modifying the risk of attack not just in terms of signals received, but in relation to the animal's own movements, such as the increased risk when emerging from cover.

Empirical support for the existence of two complementary threat detection systems in humans and other primates has come almost exclusively from studies of the visual system (e.g. Morris *et al.* 1999; Ohman 2005). For example, the cortical processing of visual images can be disrupted by a process of 'backward masking,' in which very briefly presented words or pictures are immediately followed by other scrambled or unrelated words or pictures. When such images are highly emotive (e.g. pictures of phobic stimuli such as snakes or spiders), they can cause rapid responses even though the images themselves are not registered consciously (e.g. Ohman & Soares 1994). So, threatening visual information can give rise to fearful/escape responses even in the absence of cortical processing. Moreover, two distinct neural pathways have been found to carry visual information to different brain regions at different rates. The magnocellular pathway carries low quality, particularly low spatial frequency (LSF), visual information, rapidly reaching subcortical effector sites such as the amygdala (Vuilleumier *et al.* 2003). The parvocellular pathway carries much more detailed, high spatial frequency (HSF) information to the visual cortex, albeit at a slower rate (Merigan & Maunsell 1993). Within the visual context, therefore, the magnocellular and parvocellular pathways appear to be good candidates for LeDoux's (1996) 'low road' and 'high road' to threat detection.

Despite its limited information-carrying capacity, the magnocellular pathway has been demonstrated to be capable of gathering emotionally relevant details about

stimuli, such as whether a face is showing an emotive or neutral expression (Winston *et al.* 2003; Carretié *et al.* 2007). There is also evidence that this rapidly available emotional information (e.g. registration of potential threat as a result of viewing very briefly presented images) informs and affects visual information gathering over the subsequent time period, redirecting gaze towards the affective target, particularly in anxious individuals (Mogg *et al.* 1995). This is consistent with our analysis in which the cortical decision-making system is aware of and informed by the thalamic one. Whether and to what extent the cortical decision process may convey information to the subcortical, thalamic decision-making system, however, has yet to be empirically explored.

Oliva *et al.* (2006) experimented with hybrid images, which are constructed by merging two pictures, one which has been converted to an LSF image and the other having been converted to a HSF image. The LSF image can generally be taken in 'at a glance,' whereas the HSF image requires the observer's eyes to saccade around the image to build up the picture from the thin lines of the image. Presented with a hybrid image for only 30 ms, observers reported having seen the LSF image, whereas if the image were presented for 150 ms, the HSF image was reported. In the context of this paper, if the LSF image is regarded as the signal received by the thalamic system and the HSF image can only be built by the cortical system, then the results obtained by Oliva *et al.* make sense. Further, if a hybrid image were moved (shaken) constantly so as to make it more difficult for the cortical system to perceive the thin lines and construct the HSF picture through visual saccades, we would expect that only the LSF image would be reported. This has also been observed by A. Oliva *et al.* (2008, personal communication).

The extent to which each subsystem is optimized with respect to the other (in real time, rather than through only phylogenetic effects) could be tested by loading the cortex with additional processing tasks. If the thalamus were then to make more decisions (to compensate for the slowed cortex), there would be good evidence for direct information flow between the subsystems, not only of decisions made but also of expected subsystem performance. From an experimental perspective, results may be complicated by animals having additional choices available to them when sensing danger, such as freezing, which are not included in this model. However, it may be possible, through the use of hybrid images, to experiment on humans in stressed and unstressed conditions to determine the extent of the two-way optimization between the subsystems.

This work was supported by the EPSRC. E.P. was supported by the BBSRC.

REFERENCES

- Bogacz, R. 2007 Optimal decision-making theories: linking neurobiology with behaviour. *Trends Cogn. Sci.* **11**, 18–125. (doi:10.1016/j.tics.2006.12.006)
- Bogacz, R., Brown, E., Moehlis, J., Holmes, P. & Cohen, J. D. 2006 The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced choice tasks. *Psychol. Rev.* **113**, 700–765. (doi:10.1037/0033-295X.113.4.700)

- Bogacz, R., Hu, P., Holmes, P. & Cohen, J. D. Submitted. Do humans select the speed-accuracy tradeoff maximizing reward rate?
- Brown, J. S. & Kotler, B. P. 2004 Hazardous duty pay and the foraging cost of predation. *Ecol. Lett.* **7**, 999–1014. (doi:10.1111/j.1461-0248.2004.00661.x)
- Carpenter, R. H. S. & Williams, M. L. L. 1995 Neural computation of log likelihood in control of saccadic eye movements. *Nature* **377**, 59–62. (doi:10.1038/377059a0)
- Carretié, L., Hinojosa, J. A., López-Martín, S. & Tapia, M. 2007 An electrophysiological study on the interaction between emotional content and spatial frequency of visual stimuli. *Neuropsychologia* **45**, 1187–1195. (doi:10.1016/j.neuropsychologia.2006.10.013)
- Chittka, L., Dyer, A. G., Bock, F. & Dornhaus, A. 2003 Bees trade off foraging speed for accuracy. *Nature* **424**, 388. (doi:10.1038/424388a)
- Cisek, P. & Kalaska, J. F. 2005 Neural correlates of reaching decisions in dorsal premotor cortex: specification of multiple direction choices and final selection of action. *Neuron* **45**, 801–814. (doi:10.1016/j.neuron.2005.01.027)
- Ditterich, J. 2006a Stochastic models of decisions about motion direction: behaviour and physiology. *Neural Netw.* **19**, 981–1012. (doi:10.1016/j.neunet.2006.05.042)
- Ditterich, J. 2006b Evidence for time-variant decision making. *Eur. J. Neurosci.* **24**, 3628–3641. (doi:10.1111/j.1460-9568.2006.05221.x)
- Dukas, R. (ed.) 1998 *Cognitive ecology*. Chicago, IL: University of Chicago Press.
- Egan, J. P. 1975 *Signal detection theory and ROC analysis*, pp. 16–18. New York, NY: Academic Press.
- Gold, J. & Shadlen, M. 2002 Banburismus and the brain: decoding the relationship between sensory stimuli, decisions, and reward. *Neuron* **36**, 299–308. (doi:10.1016/S0896-6273(02)00971-6)
- Houston, A. I. & McNamara, J. M. 1986 Evaluating the selection pressure on foraging decisions. In *Relevance of models and theories in ethology* (eds R. Campan & R. Zayan), pp. 61–75. Toulouse, France: Private Press.
- Houston, A. I. & McNamara, J. M. 1999 *Models of adaptive behaviour*. Cambridge, UK: Cambridge University Press.
- Laming, D. R. J. 1968 *Information theory of choice reaction times*. London, UK: Academic Press.
- LeDoux, J. E. 1996 *The emotional brain*. New York, NY: Simon & Schuster.
- Lima, S. L. 1998 Stress and decision making under the risk of predation: recent developments from behavioral, reproductive, and ecological perspectives. *Adv. study Behav.* **27**, 251–290. (doi:10.1016/S0065-3454(08)60366-6)
- MacLean, P. D. 1982 On the origin and progressive evolution of the triune brain. In *Primate brain evolution, methods and concepts* (eds E. Armstrong & D. Falk), pp. 291–316. New York, NY: Plenum Press.
- Maddox, W. T. 2002 Toward a unified theory of decision criterion learning in perceptual categorization. *J. Exp. Anal. Behav.* **78**, 567–595. (doi:10.1901/jeab.2002.78-567)
- McHaffie, J. G., Stanford, T. R., Stein, B. E., Coizet, V. & Redgrave, P. 2005 Subcortical loops through the basal ganglia. *Trends Neurosci.* **28**, 401–407. (doi:10.1016/j.tins.2005.06.006)
- Merigan, W. H. & Maunsell, J. H. 1993 How parallel are the primate visual pathways? *Annu. Rev. Neurosci.* **16**, 369–402. (doi:10.1146/annurev.ne.16.030193.002101)
- Mobbs, D., Petrovic, P., Marchant, J. L., Hassabis, D., Weiskopf, N., Seymour, B., Dolan, R. J. & Frith, C. D. 2007 When fear is near: threat imminence elicits prefrontal-periaqueductal gray shifts in humans. *Science* **317**, 1079–1083. (doi:10.1126/science.1144298)
- Mogg, K., Bradley, B. P. & Williams, R. 1995 Attentional bias in anxiety and depression: the role of awareness. *Br. J. Clin. Psychol.* **34**, 17–36.
- Morris, J. S., Ohman, A. & Dolan, R. J. 1999 A subcortical pathway to the right amygdala. *Proc. Natl Acad. Sci. USA* **96**, 1680–1685. (doi:10.1073/pnas.96.4.1680)
- Ohman, A. 2005 The role of the amygdala in human fear: automatic detection of threat. *Psychoneuroendocrinology* **30**, 953–958. (doi:10.1016/j.psyneuen.2005.03.019)
- Ohman, A. & Soares, J. J. F. 1994 'Unconscious anxiety': phobic responses to masked stimuli. *J. Abnorm. Psychol.* **103**, 231–240. (doi:10.1037/0021-843X.103.2.231)
- Oliva, A., Torralba, A. & Schyns, P. G. 2006 Hybrid images. *ACM Trans. Graph.* **25**, 527–532. (doi:10.1145/1141911.1141919)
- Ratcliff, R. 1978 A theory of memory retrieval. *Psychol. Rev.* **83**, 59–108. (doi:10.1037/0033-295X.85.2.59)
- Ratcliff, R. & Smith, P. L. 2004 Comparison of sequential sampling models for two-choice reaction time. *Psychol. Rev.* **111**, 333–367. (doi:10.1037/0033-295X.111.2.333)
- Real, L. A. 1992 Information-processing and the evolutionary ecology of cognitive architecture. *Am. Nat.* **140**, S108–S145. (doi:10.1086/285399)
- Real, L. A. 1993 Toward a cognitive ecology. *Trends Ecol. Evol.* **8**, 413–417. (doi:10.1016/0169-5347(93)90044-P)
- Roitman, J. D. & Shadlen, M. N. 2002 Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J. Neurosci.* **22**, 9475–9489.
- Schall, J. D. 2001 Neural basis of deciding, choosing and acting. *Nat. Rev. Neurosci.* **2**, 33–42. (doi:10.1038/35049054)
- 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. (doi:10.1016/j.tins.2004.01.006)
- Smith, J. D., Schull, J., Strote, J., McGee, K., Egnor, R. & Erb, L. 1995 The uncertain response of the bottlenosed dolphin (*Tursiops truncatus*). *J. Exp. Psychol. Gen.* **124**, 391–408. (doi:10.1037/0096-3445.124.4.391)
- Sole, L. M., Shettleworth, S. J. & Bennett, P. J. 2003 Uncertainty in pigeons. *Psychon. Bull. Rev.* **10**, 738–745.
- Usher, M. & McClelland, J. L. 2001 The time course of perceptual choice: the leaky, competing accumulator model. *Psychol. Rev.* **108**, 550–592. (doi:10.1037/0033-295X.108.3.550)
- Vuilleumier, P., Armony, J. L., Driver, J. & Dolan, R. J. 2003 Distinct spatial frequency sensitivities for processing faces and emotional expressions. *Nat. Neurosci.* **6**, 624–631. (doi:10.1038/nn1057)
- Wald, A. 1945 Sequential tests of statistical hypotheses. *Ann. Math. Stat.* **16**, 117–186. (doi:10.1214/aoms/1177731118)
- Winston, J. S., Villeumier, P. & Dolan, R. J. 2003 Effects of low-spatial frequency components of fearful faces on fusiform cortex activity. *Curr. Biol.* **13**, 1824–1829. (doi:10.1016/j.cub.2003.09.038)
- Wong, K. F. & Wang, X. J. 2006 A recurrent network mechanism of time integration in perceptual decisions. *J. Neurosci.* **26**, 1314–1328. (doi:10.1523/JNEUROSCI.3733-05.2006)
- Yang, T. & Shadlen, M. N. 2007 Probabilistic reasoning by neurons. *Nature* **447**, 1075–1080. (doi:10.1038/nature05852)
- Zald, D. H. 2003 The human amygdala and the emotional evaluation of sensory stimuli. *Brain Res. Rev.* **41**, 88–123. (doi:10.1016/S0165-0173(02)00248-5)