

A Bayesian Model of the Rule Association Task and Effect of Priming

S. Bengtsson and W. Penny

Wellcome Trust Centre for Neuroimaging,
University College, London WC1N 3BG, UK.

April 30, 2013

Abstract

This paper describes a normative model of the Rule Association task based on Bayesian inference and formulated in discrete time. We also show how it reduces to a standard Drift Diffusion Model in the continuum limit. This relationship then provides semi-analytic formulae for the mean reaction time and error rate. The model is applied to behavioral data from an experiment where we used priming to manipulate subjects ability to perform the task.

1 Introduction

Bayesian inference has proven useful in understanding brain and behaviour on many levels [5], from sensory perception [6] to motor learning [9] and social interaction [13], whereas no study has yet been conducted on the effect of self-associations on behaviour. This paper derives a normative model of the rule association task based on such principles.

Recently Yu et al. [14] have described a normative model of the Eriksen Flanker task, based on Bayesian inference. In later work [10] they also provide a connection to Drift Diffusion Models (DDMs) from which they derive semi-analytic formulae for reaction times and error rates. The work in this paper is inspired by their approach and applies a similar methodology to a rule association task.

2 Theory

In the rule association task [4], as depicted in Figure 1, there are two types of trials (i) univalent trials in which the mapping from a stimulus to response is

fixed and (ii) bivalent trials in which the identity of a contextual cue reverses a stimulus-response mapping. The responses are either left or right button presses.

2.1 Discrete-Time Model

We define generative models in which stimuli s and context c produce observations of neural activity related to the stimulus, $x_s(t)$, and to the context, $x_c(t)$, at time t .

For the Rule task c is taken to be the identity of the cue defining the rule, and s the identity of the stimulus. The corresponding neural activities are then given by

$$\begin{aligned} p(x_s(t)|s) &= \mathbf{N}(x_s(t); \mu_s, \sigma_s^2) \\ p(x_c(t)|c) &= \mathbf{N}(x_c(t); \mu_c, \sigma_c^2) \end{aligned} \quad (1)$$

where we take $\mu_s = -1$ for $s = T$ (Tree), $\mu_s = 1$ for $s = B$ (Butterfly) and $\mu_c = -1$ for $c = C$ (Cross) and $\mu_c = 1$ for $c = S$ (Star). Thus, neural activity representing T and B has mean of -1 and +1 respectively, and the representational acuity of this activity is determined by the noise deviation parameter, σ_s . Similarly, σ_c determines the accuracy of the representations for C and S. This is similar to the generative model for centre and flanker stimuli in Yu et al.'s [14] model of the Eriksen Flanker task. Here we allow stimulus and context to have different levels of noise.

We now derive the recognition model, that is, given neural activity x how do we decide what stimulus is present and which response to make? We denote $x(t)$ as the vector containing both observations $x_s(t)$, $x_c(t)$, and X_T as the vector containing all $x(t)$'s up to time T .

For the rule task we can write down a recursive formula for the joint probability of X_t and the response, r , being left (L)

$$\begin{aligned} p(r = L, X_t) &= [p(x_c(t)|c = C)p(x_s(t)|s = T) \\ &+ p(x_c(t)|c = S)p(x_s(t)|s = B)]p(r = L|X_{t-1}) \end{aligned} \quad (2)$$

This encodes the rule that one should press left if the context is a cross (C) and the stimulus a tree (T), or if the context is a star (S) and the stimulus a butterfly (B). Similarly, one can define an update for the joint probability of the observation and a right button response

$$\begin{aligned} p(r = R, X_t) &= [p(x_c(t)|c = C)p(x_s(t)|s = B) \\ &+ p(x_c(t)|c = S)p(x_s(t)|s = T)]p(r = R|X_{t-1}) \end{aligned} \quad (3)$$

One can then compute the probability of a left button response as

$$p(r = L|X_t) = \frac{p(r = L, X_t)}{p(r = L, X_t) + p(r = R, X_t)} \quad (4)$$

A decision is then made to press the left button if this probability exceeds a threshold, β , to press the right button if it drops below $1-\beta$, or to keep observing if neither condition is met. Figure 2 shows how the probability, $p(r = L|X_t)$ evolves over time. Together, the within-trial dynamics defined by the above equations and the threshold β determine the reaction time of the normative subject.

2.1.1 Univalent task

For the univalent task the stimulus-response mapping is fixed throughout the course of the experiment. Because of this we envisage that the stimulus itself is sufficient to perform the task and that post-stimulus integration of stimulus with cue information is unnecessary. In modelling terms we can therefore ignore x_c . The joint probability of X_t and a left or right responses are therefore

$$\begin{aligned} p(r = L, X_t) &= p(x_s(t)|s = H)p(r = L|X_{t-1}) \\ p(r = R, X_t) &= p(x_s(t)|s = C)p(r = L|X_{t-1}) \end{aligned} \quad (5)$$

where $s = H$ and $s = C$ denote house and car stimuli. One can then form the posterior probabilities from equation 4. Given uniform prior probabilities $p(r = L|X_0) = p(r = R|X_0) = 0.5$, a decision based on the posterior probability is identical to one based on the likelihood, or log-likelihood ratio

$$I_t = \log \left(\frac{p(X_t|r = L)}{p(X_t|r = R)} \right) \quad (6)$$

This can be accumulated sequentially as

$$I_t = I_{t-1} + \delta I_t \quad (7)$$

where

$$\delta I_t = \log \left(\frac{p(x_t|r = L)}{p(x_t|r = R)} \right) \quad (8)$$

2.2 Continuous-Time Model

The decision process for the univalent task is a variant of the Sequential Probability Ratio Test (Wald, 1947) which is known to optimal for Two-Alternative Forced Choice (2AFC) tasks. SPRTs can also be formulated in continuous time and are known as Drift-Diffusion Models (DDMs) [2]. A benefit of the DDM approach is that, for simple models, there are analytic expressions for the reaction time and error rates. This obviates the need for implementing computationally intensive stochastic simulations of ideal subject responses, and therefore makes it much simpler to fit these models to data.

2.2.1 Univalent Rule Task

This section shows how the DDM is derived as a continuum limit of the discrete time evidence accumulation process implemented for the univalent rule task. The DDM takes the form of the Stochastic Differential Equation (SDE)

$$dI = mdt + Ddw \quad (9)$$

with drift parameter m and diffusion parameter D , where dw is a standard Wiener process (Gardner, 1983; Tuckwell, 1995). The expected correct rate, CR, and expected decision time, DT, are then given by analytic expressions [2]

$$\begin{aligned} CR &= \frac{1}{1 + \exp(-2mz/D^2)} \\ DT &= \frac{z}{m} \tanh\left(\frac{mz}{D^2}\right) \\ z &= \log\left(\frac{\beta}{1-\beta}\right) \end{aligned} \quad (10)$$

Here, z is the decision threshold in the units of a log odds ratio. The reaction time is then given by

$$RT = DT + T_0 \quad (11)$$

where T_0 accounts for the time taken by sensorimotor processes that are not part of the decision process.

To make the connection with discrete time models we note that the above SDE can be numerically integrated using the Euler-Maruyama (EM) method (Higdon, 2001) to give a discrete-time update of the form

$$\Delta I_n = m\Delta t + D(w(t_n) - w(t_{n-1})) \quad (12)$$

where the noise increment $w(t_n) - w(t_{n-1})$ is drawn from a Gaussian distribution with mean 0 and variance Δt (that is, a Wiener process). This update can then be written as

$$\Delta I_n = m\Delta t + D\sqrt{\Delta t}e_n \quad (13)$$

where e_n is now a standard Gaussian variate. Hence, the central moments are given by

$$\begin{aligned} E(\Delta I_n) &= m\Delta t \\ Var(\Delta I_n) &= D^2\Delta t \end{aligned} \quad (14)$$

These equations can be rearranged to give m and D in terms of the moments.

In the discrete time formulation, evidence is accumulated as

$$\Delta I_n = \log \frac{p(x|r_L)}{p(x|r_R)} \quad (15)$$

If each of these densities is a single Gaussian with mean μ_1 for left, μ_2 for right, and with common variance σ^2 then the moments are

$$\begin{aligned} E(\Delta I_n) &= \frac{(\mu_1 - \mu_2)^2}{2\sigma^2} \\ \text{Var}(\Delta I_n) &= \frac{(\mu_1 - \mu_2)^2}{\sigma^2} \end{aligned} \quad (16)$$

Equating these with those in equation 16 gives the following expressions

$$\begin{aligned} m &= \frac{(\mu_1 - \mu_2)^2}{2\sigma^2 \Delta t} \\ D &= \frac{(\mu_1 - \mu_2)}{\sigma \sqrt{\Delta t}} \end{aligned} \quad (17)$$

These values can then be used in equation 10 to compute CR and DT. The expression for D differs from equation A23 in Bogacz et al. (2006) which is incorrect. We have also validated these expressions using Monte-Carlo simulations and found it necessary to use a very small Δt so that the Euler approximation was valid. This is well known in the stochastic processing literature (Higdon, 2001) and has been examined in the context of DDMs by Brown et al. (2006) who also compares different integration methods. In what follows in this paper we use the value $\Delta t = 1\text{ms}$.

The expressions for m and D in equation 17 mean that $m/D^2 = 1/2$. Hence, for accumulation of information from a single source as in the univalent rule task, the expressions for correct rate and decision time reduce to

$$\begin{aligned} CR &= \beta \\ DT &= \frac{z}{m} (2\beta - 1) \\ z &= \log \left(\frac{\beta}{1 - \beta} \right) \end{aligned} \quad (18)$$

This is intuitively satisfying because the correct rate in this simple case is exactly equal to the probability threshold β . We now derive similar expressions for the bivalent rule task.

2.2.2 Bivalent Rule Task

Again, assuming equal priors, a decision based on posterior probabilities is equivalent to one based on likelihood ratios. For the bivalent rule task evidence is accumulated according to

$$\Delta I_n(x_c, x_s) = \log \frac{p(r = L, X_t)}{p(r = R, X_t)} \quad (19)$$

By using expressions for the joint densities in equations 2 and 3 this can be re-arranged as

$$\Delta I_n(x_c, x_s) = \frac{\exp(-\frac{1}{2}z_{c1}^2 - \frac{1}{2}z_{s1}^2) + \exp(-\frac{1}{2}z_{c2}^2 - \frac{1}{2}z_{s2}^2)}{\exp(-\frac{1}{2}z_{c1}^2 - \frac{1}{2}z_{s2}^2) + \exp(-\frac{1}{2}z_{c2}^2 - \frac{1}{2}z_{s1}^2)} \quad (20)$$

where

$$\begin{aligned}
z_{c1} &= (x_c - \mu_c)/\sigma_c \\
z_{c2} &= (x_c - \mu_c)/\sigma_c \\
z_{s1} &= (x_s - \mu_s)/\sigma_s \\
z_{s2} &= (x_s - \mu_s)/\sigma_s
\end{aligned} \tag{21}$$

where c and s index the context (Cross or Star) and stimulus (Tree or Butterfly) on that trial. We then have

$$\begin{aligned}
g_m(\sigma_c, \sigma_s) &= E(\Delta I_n) \\
&= \int \int \Delta I_n(x_c, x_s) p(x_c) p(x_s) dx_s dx_c
\end{aligned} \tag{22}$$

where $p(x_c)$ and $p(x_s)$ are defined in equation 1. Unfortunately, we know of no analytic solution to this integral, but as it is only two dimensional it can be quickly evaluated using numerical quadrature. That is

$$g_m(\sigma_c, \sigma_s) = \sum_{i=1}^{N_m} \sum_{j=1}^{N_m} \Delta I_n(x_c(i), x_s(j)) p(x_c(i)) p(x_s(j)) \tag{23}$$

The variance term can then be evaluated in a similar manner

$$\begin{aligned}
g_v(\sigma_c, \sigma_s) &= Var(\Delta I_n) \\
&= \sum_{i=1}^{N_m} \sum_{j=1}^{N_m} (\Delta I_n(x_c(i), x_s(j)) - g_m(\sigma_c, \sigma_s))^2 p(x_c(i)) p(x_s(j))
\end{aligned} \tag{24}$$

It turns out that both integrals can be accurately estimated using only $N_m = 10$ points in each dimension (100 points in all), where x_s, x_c are evenly spaced between plus and minus three σ_c, σ_s . This computation takes about 10ms on a standard personal computer. We have validated this approach by comparing with Monte-Carlo simulations in which x_s and x_c are drawn from equation 1 and ΔI_n is evaluated using equation 20. This was repeated $S = 10,000$ times and all samples were used to compute the mean and variance. Figure 3 compares the numerical quadrature evaluations against Monte-Carlo estimates, showing good agreement over a broad range of values. These values include the estimates found in model fitting (see below). Agreement is similarly good for the variance term. In what follows, we therefore compute the moments using the quadrature approach.

As there are two possible contexts and two possible stimuli, we have four different trial types within the bivalent task, and must therefore evaluate the moments for each trial type. However, because of the symmetry of the task, g_v and g_m are the same for all trial types (actually g_m is flipped for right as opposed to left responses but then evidence is accumulated in the other direction).

The parameters of the equivalent SDE are then given by

$$\begin{aligned} m &= \frac{g_m(\sigma_c, \sigma_s)}{\Delta t} \\ D^2 &= \frac{g_v(\sigma_c, \sigma_s)}{\Delta t} \end{aligned} \tag{25}$$

These values can then be plugged into equation 10 to give predictions of CR and DT. Importantly, we have found that the ratio

$$r = \frac{g_m}{g_v} \tag{26}$$

has a very weak dependence on σ_s and σ_c . Specifically, we varied σ_s and σ_c between 1 and 10, with 50 linearly-spaced values per dimension, and evaluated r at each point. These values correspond to the range observed empirically (see below). The minimum $r = 0.474$ was found at $\sigma_s = \sigma_c = 1$ and the maximum $r = 0.508$ at $\sigma_s = \sigma_c = 10$. The mean and standard deviation were $E[r] = 0.499$ and $\text{std}(0.008)$. For a decision threshold of $z = 2$ the minimum and maximum r values translate to correct rates of 87.0% and 88.4%, a minimal change. In what follows we therefore set $r = 0.5$. This means that (i) as for the univalent case, the correct rate is determined solely by the decision threshold, $CR = \beta$ and (ii) there is no need to compute g_v as it only enters the DT equation as g_m/g_v .

2.2.3 Decision thresholds

Increasing the threshold β leads to slower, more accurate decisions and decreasing it to faster, less accurate ones. SPRTs and DDMs can optimize this speed-accuracy trade-off, by setting the threshold β appropriately. Gold and Shadlen (2002) have proposed that participants select the threshold, β , that maximises the reward rate (RR), defined as the average number of rewards per unit time. In the context of our experiment this is simply the number of correct responses per unit time.

According to Bogacz et al [3], the speed accuracy trade-off is implemented by neural circuits in association and pre-motor areas. Further, rather than a threshold being changed, it is the initial value of the accumulation point which is thought to be modulated. DDMs are agnostic as to whether the initial or final point (threshold) is changed, because both produce identical behaviour. Bogacz et al [3] suggest that the change in initial value is mediated by increasing the baseline firing rate of cortical integrator neurons.

2.3 Model Fitting

The DDM for the univalent rule task comprises the following parameters; T_0 , σ_s and z . The DDM for the bivalent task has the additional parameter σ_c . We estimated model parameters using data from all subjects, rather than subject

by subject which is a common approach in fitting DDMMs to behavioural data [10, 12].

Model fitting was carried out as follows. We used adjusted RT data (with subject means subtracted and overall group means added as described in the results section below) so that we could use a single T_0 parameter for all subjects. For both univalent and bivalent data, decision thresholds z_{uni} and z_{bi} were computed for each session and subject directly from the CR values (see equation 18).

We then constructed a model for predicting the univalent and bivalent reaction times from the known values z_{uni} , z_{bi} and unknown values T_0 , σ_s and $\sigma_c(j)$ where j indexes the j th session of the bivalent data.

These unknown values were parameterised as follows

$$\begin{aligned} T_0 &= \exp(\theta_T) \\ \sigma_s &= \exp(\theta_s) \\ \sigma_c(j) &= \frac{\sigma_s}{1 + w_j} \\ w_j &= \exp[\theta_c(j)] \end{aligned} \tag{27}$$

where the last two lines incorporate two assumptions. First, the contextual noise variance is constrained to be smaller than the stimulus noise variance. We also consider a model (see Results) in which the converse is true. Because model RTs are symmetric in σ_s and σ_c , enforcing one to be larger than the other helps to produce an identifiable model.

Second, because our larger goal is to study the effect of priming on decision making we hypothesized that σ_c varies between primes and over sessions. Obviously, at least one of the parameters has to change to accommodate the between-session difference in error rates (see below), and it seems that the most likely candidate is contextual rather than stimulus noise. This is because priming did not affect the univalent task (and therefore not σ_s).

This rather constrained model gives us five parameters to estimate from 90 data points. We also consider a more constrained version, with three parameters, in which $\theta_c(j)$ does not vary over session ie. $w_j = w$. We refer to this as the null model and the less constrained version as the alternative model. Comparing the two will allow us to infer whether σ_c changes over sessions.

The model is then fitted using a Bayesian parameter estimation scheme, known as the Variational Laplace algorithm [7]. The algorithm is Bayesian, in that priors are placed over the parameters to be estimated, which in our case are the latent variables θ_T , θ_s and $\theta_c(j)$. We used the following Gaussian priors

$$\begin{aligned} p(\theta_T) &= \mathbf{N}(\mu_T, v_T^2) \\ p(\theta_s) &= \mathbf{N}(\mu_s, v_s^2) \\ p(\theta_c(j)) &= \mathbf{N}(\mu_c, v_c^2) \end{aligned} \tag{28}$$

The implicit prior distributions over T_0 , σ_s and w_j therefore have a log-normal form. We used uninformative priors over σ_s and w_j obtained using the values

$\mu_s = 0.69, v_s = 0.5$ and $\mu_c = 0, v_c = 1$. To obtain realistic values of the nondecision time, T_0 , we found it necessary to use a rather informative prior. This constrained T_0 to have a prior mean of 500ms, which is similar to values from other DDM fits [12, 2]. The prior was obtained using the values $\mu_T = -0.69$ and $v_T = 1/64$.

The VL algorithm then estimates the parameters of approximate posterior densities over θ using a pseudo-Newton like algorithm [7, 11]. Moreover, it also produces an estimate of the model evidence which provides a quantitative measure for comparing models based on Bayes factors [8]. Evidence in favour of model 1 versus model 2 is assessed using the Bayes Factor

$$BF(1, 2) = \frac{p(y|m_1)}{p(y|m_2)} \quad (29)$$

and the posterior probability of model 2 is given by

$$p(m_2|y) = \frac{1}{1 + BF(1, 2)} \quad (30)$$

Null and alternative models can be compared using these posterior probabilities as will be shown in the results section.

3 Experiment

15 native English speaking volunteers (aged 24.7 ± 4.1 yrs; 8 females) took part in the study. The participants all gave written informed consent, and the study was approved by the joint ethics committee of the Institute of Neurology and University College London Hospital, London, UK. Each subject took part in a total of six experimental sessions, and each session comprised a priming part followed by a rule association part.

Priming refers to the passive, subtle, and unobtrusive activation of relevant mental representations by external, environmental stimuli, such that people are not and usually do not become aware of the influence exerted by those stimuli [1].

In the lab, priming can be undertaken with the scrambled sentence task. In our study, each scrambled sentence consisted of six words and participants judged whether or not it could be made into a grammatically coherent sentence by using five of the six words. In each session, 70% of the sentences had words that were synonyms for either Clever or Stupid, and 30% of the sentences were neutral. The effect of priming was evaluated on the Rule Association task.

Eight sentences were presented followed by a sequence of 50 rule trials. This constitutes a session and there were three consecutive sessions for each prime (clever and stupid). The order of the two conditions (clever and stupid) was pseudo-randomized between participants. Prior to data collection, participants practised the rule task for 80 trials, and the language task for 20 sentences, with all the sentences being of a neutral character.

4 Results

4.1 Univalent Task

Overall there was no difference in the error rate on the univalent task when primed ‘stupid’ ($6.9 \pm 1.8\%$) versus ‘clever’ ($6.0 \pm 0.6\%$) (paired t-test, $p < 0.6$). Similarly, there was no difference in reaction time for ‘stupid’ ($604 \pm 32\text{ms}$) versus ‘clever’ ($643 \pm 45\text{ms}$) (paired t-test, $p < 0.12$). There was a tendency for clever-primed subjects to get faster, as mean RTs in sessions 1, 2 and 3 were 877, 849 and 816ms. However, neither of these effects were significant (one-way ANOVA, $p = 0.29$ for stupid and $p = 0.11$ for clever).

Behavioural data were therefore collapsed over the two prime types. RTs and CRs were extracted over all subjects and sessions giving rise to 45 data points from three sessions and fifteen subjects. The overall mean RT was $628 \pm 38\text{ms}$ and the mean CR was $94.0 \pm 1\%$.

Because univalent and bivalent trials were mixed together in the same session, reaction times may also be affected by recency effects (Cho et al), for example, whether the previous trial was also a univalent (‘stay’) or bivalent (‘switch’) trial. However, the difference in reaction time for switch minus stay trials was not significant ($15 \pm 15\text{ms}$, paired t-test, $p = 0.30$).

4.2 Bivalent Task

Overall, the error rate on the bivalent task was significantly higher when subjects had been primed ‘stupid’ (12.2 ± 2.0), as compared to ‘clever’ (8.1 ± 1.2) ($p < 0.001$). The error rates when participants were primed stupid were significantly different between session (one-way ANOVA, $p = 6 \times 10^{-6}$). The error rates were 8, 12 and 18 percent for sessions 1 to 3. The error rates when participants were primed clever were not significantly different between session (one-way ANOVA, $p = 0.64$). The error rates were 8, 9 and 8 percent for sessions 1 to 3. Thus, we observed that when stupid associations are evoked participants’ performance becomes increasingly worse. We refer to this as a confirmation bias. Boxplots of the error rate effects are shown in Figure 4. The confirmation bias effect remained highly significant ($p = 1 \times 10^{-5}$) when the outlying participant (see circles in top row of Figure 4) was removed from the analysis.

Overall, when subjects were primed stupid as opposed to clever they were significantly faster (822 versus 874ms, paired t-test, $p = 0.03$). There was a tendency for stupid-primed subjects to slow down with session and clever-primed subjects to speed up, but neither of these effects were significant (one-way ANOVA, $p = 0.29$ for stupid and $p = 0.11$ for clever).

The difference in reaction time for switch minus stay trials was significant ($89 \pm 21\text{ms}$, paired t-test, $p = 10^{-4}$). Given that 29% trials were switch trials this therefore has the effect of increasing the mean reaction time in each session by $0.29 \times 89 = 26\text{ms}$. We therefore subtracted this amount from the mean reaction time of each subject and session, prior to model fitting.

4.3 Model Fitting

We now present results on fitting the DDM rule task models to the behavioural data presented in the previous sections. We first report perhaps the most important inference which is whether σ_c varies over sessions. This is assessed using Bayes factors. For the data obtained when subjects were primed stupid, the log Bayes factor favouring the alternative model over the null is 17.6. This corresponds to a posterior belief in the alternative model approaching unity ($p(alt|y) = 1 - 10^{-8}$). For the clever data the log Bayes factor is 0.96 and the posterior belief is $p(alt|y) = 0.72$. Hence, we infer that σ_c varies over session for the stupid data but not for the clever. This reflects the fact that error rates changed for the stupid data but not the clever.

For the stupid data, the parameters were estimated to be $T_0 = 520\text{ms}$ and $\sigma_s = 7.4$. The between-session weighting factors that set σ_c were estimated to be $w_1 = 4.57$, $w_2 = 3.05$ and $w_3 = 2.59$. This corresponds to σ_c values of 1.63, 2.44 and 2.87. That is a 76% increase in σ_c from session 1 to session 3. This latter result is robust over a range of different settings for the prior distributions. Figure 5 shows mean reaction times for data and model predictions on the univalent and bivalent tasks (stupid prime).

Figure 6 shows the how reaction time on the bivalent task is predicted to vary as a function of σ_c and z . The values were computed assuming the estimates $T_0 = 520\text{ms}$ and $\sigma_s = 7.4$. Also, marked are the estimated z and σ_c values for sessions 1, 2 and 3 when subjects were primed stupid.

We also fitted models in which the contextual noise was constrained to be larger than the stimulus noise, but there was significantly less evidence for these models. The posterior probability of the $\sigma_s > \sigma_c$ model was 0.9999. The $\sigma_s < \sigma_c$ model could only provide a good fit to the data by using an apriori unlikely large value of T_0 . Importantly, however, even for the $\sigma_s < \sigma_c$ models our main finding was upheld. That is, that σ_c varies over session for the stupid data but not for the clever.

5 Discussion

We have described a simple Bayes optimal inferential process for the Rule Association task framed in discrete time. In the continuum limit of the time step reducing to zero the process becomes equivalent to a Drift Diffusion Model. This equivalence allows us to derive simple deterministic equations for predicting behavioural data. This simplifies model fitting, as for each setting of model parameters, we do not have to resort to repeated simulation of a stochastic process.

For the univalent Rule association trials the error rate is determined solely by the decision threshold, and for the bivalent trials this is approximately the case. In the bivalent case, the dependence on other parameters such as stimulus and context noise level is very weak, and can be neglected for practical purposes.

Our results show that stupid priming significantly increases the noise level of

contextual representations. The increased noise level on its own would lead to vastly increased reaction times, as indicated by Figure 6. However, we envisage that automatic mechanisms that control the speed-accuracy trade-off [3] then act so as to bring reaction times back into a reasonable range. The consequence of this is the reduction in accuracy seen in later sessions. This provides a mechanistic explanation of the confirmation bias and reaction time data.

References

- [1] J Bargh, M Chen, and L Burrows. Automaticity of social behaviour: Direct effects of trait constructs and stereotype activation on action. *Journal of Personality and Social Psychology*, 71:230–244, 1996.
- [2] R Bogacz, E Brown, J Moehlis, P Holmes, and J Cohen. The physics of optimal decision making: a formal analysis of models and performance in two-alternative forced-choice tasks. *Psychological Review.*, 113(4):700–765, 2006.
- [3] R Bogacz, E Wagenmakers, B Forstmann, and S Nieuwenhuis. The neural basis of the speed-accuracy tradeoff. *Trends in Neuroscience*, 33(1), 2010.
- [4] E Crone, S Donohue, R Honomichl, C Wendelken, and S Bunge. Brain regions mediating flexible rule use during development. *Journal of Neuroscience*, 26:11239–11247, 2006.
- [5] K Doya, S Ishii, A Pouget, and R Rao, editors. *Bayesian Brain: Probabilistic Approaches to Neural Coding*. MIT Press, 2007.
- [6] M Ernst and M Banks. Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*, 415(6870):429–33, 2002.
- [7] K. Friston, J. Mattout, N. Trujillo-Barreto, J. Ashburner, and W. Penny. Variational free energy and the Laplace approximation. *Neuroimage*, 34(1):220–234, 2007.
- [8] R.E. Kass and A.E. Raftery. Bayes factors. *Journal of the American Statistical Association*, 90:773–795, 1995.
- [9] K Kording and D Wolpert. Bayesian integration in sensorimotor learning. *Nature*, 427:244–247, 2004.
- [10] Y Liu, A Yu, and P Holmes. Dynamical analysis of bayesian inference models for the eriksen task. *Neural Computation*, 21:1520–53, 2009.
- [11] W. H. Press, S.A. Teukolsky, W.T. Vetterling, and B.V.P. Flannery. *Numerical Recipes in C*. Cambridge, 1992.
- [12] R Ratcliff and P L Smith. A comparison of sequential sampling models for two-choice reaction time. *Psychological Review*, 111(2):333–67, 2004.

- [13] W Yoshida, R Dolan, and K Friston. Game theory of mind. *PLOS-CB*, 4(12):e1000254, 2008.
- [14] A J Yu, P Dayan, and J D Cohen. Dynamics of attentional selection under conflict: toward a rational Bayesian account. *J Exp Psychol Hum Percept Perform*, 35(3):700–17, 2009.

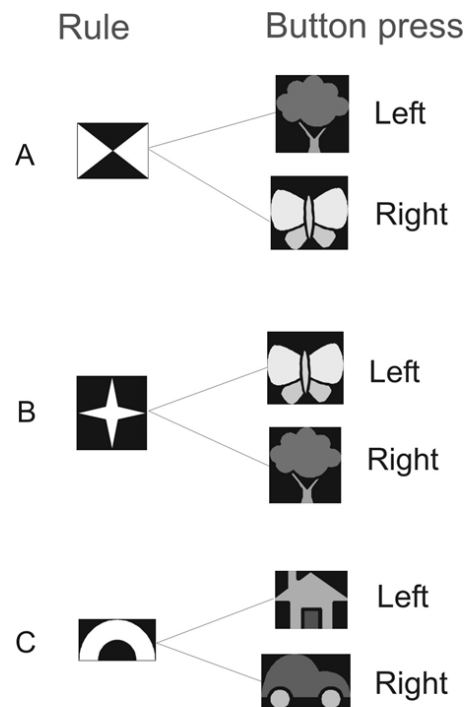


Figure 1: The experiment consisted of two rule types: A and B illustrate the bivalent rule, C illustrates the univalent rule. Participants viewed the rule cue for 1 s. After a 0.5 s delay the target stimulus was presented for 2.5 s. The response was either a left or a right button press, depending on the relevant mapping that had been previously learnt.

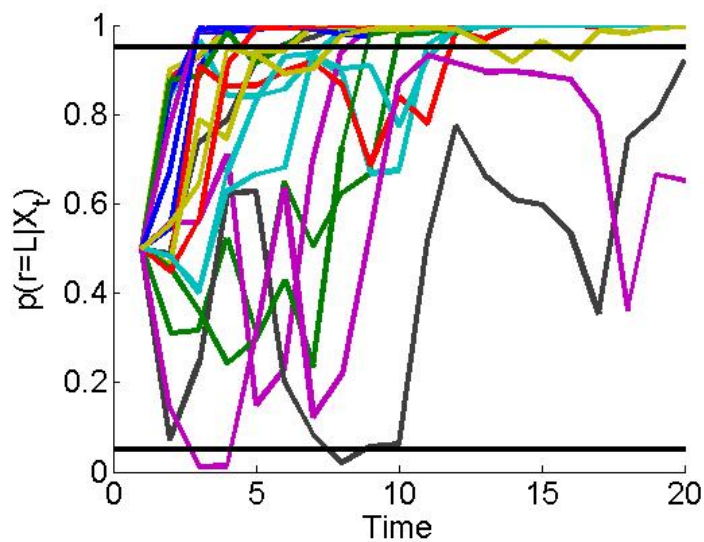


Figure 2: Simulated within-trial dynamics, governed by equations 1 to 4, for twenty trials of the rule task (each differently coloured). The true context and stimulus were $c = S$ and $s = B$, thus left is the correct response. Each trial lasts for 20 time steps and used parameters $\sigma_s = \sigma_c = 1$. The horizontal lines denote the decision thresholds β and $1 - \beta$. The (first) intersection of each trajectory with the decision threshold indicates the reaction time for each trial. There are two error trials (violet and black curves) out of twenty.




Figure 3: Monte-Carlo (blue curves) versus numerical quadrature (red curves) evaluation of $g_m(\sigma_c, \sigma_s)$ for the bivalent rule task. These estimates are plotted as a function of σ_c for various values of σ_s : 1.00 (top left), 1.84 (top right), 3.39 (bottom left) and 6.25 (bottom right).

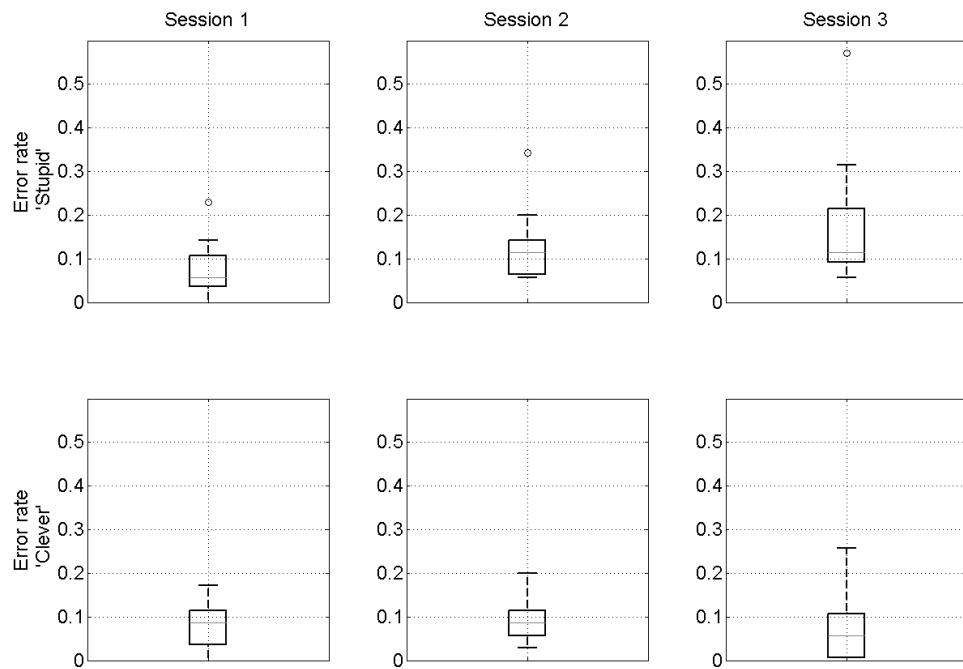


Figure 4: Boxplots of error rates over participants for 'stupid' prime (top row) and 'clever' prime (bottom row). On each box, the dotted line indicates the median, the edges of the box are the 25th and 75th percentiles, the whiskers extend to the most extreme data points considered not to be outliers, and the outliers are plotted individually as circles. The outlying data points in the top row are all from participant 10.

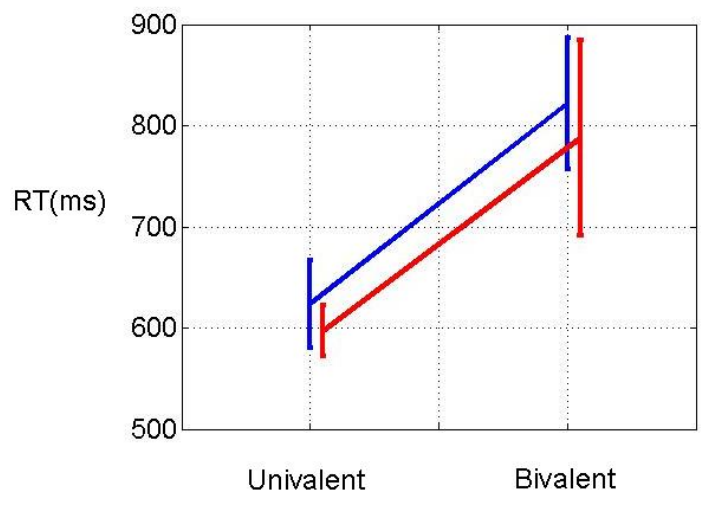


Figure 5: Mean reaction times on univalent and bivalent tasks (blue line) and model predictions (red line). The red line has been shifted to the right to make the differences more visible. Vertical bars indicate standard errors (blue line) and standard deviations (red line).

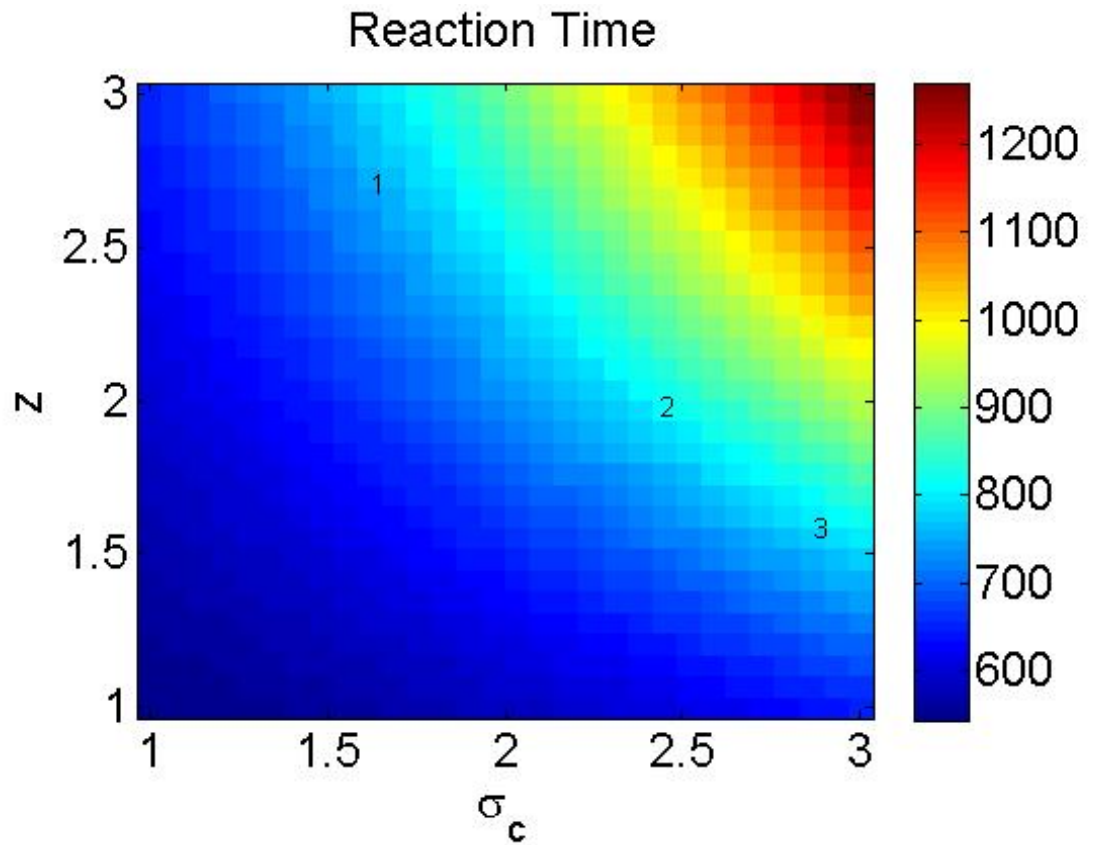


Figure 6: Model predictions of reaction time on bivalent rule task as a function of σ_c and z for $\sigma_s = 7.4$ and $T_0 = 520\text{ms}$. The symbols 1, 2 and 3 mark estimated values of z and σ_c for sessions 1, 2 and 3 when subjects were primed stupid.