SEEXC: A model of response time in skill acquisition

Andrew Heathcote and Scott Brown

School of Behavioural Sciences,

The University Of Newcastle, Australia.

email: andrew.heathcote@newcastle.edu.au

Abstract

We outline SEEXC, a neural network model of choice response time (RT) based on leaky competitive integrators. SEEXC is different from extant neutral network models in that it incorporates the effects of practice by modulating recurrent self-connection weights. For simplified versions of this model, we provide analytic and numeric results concerning RTs and the relationship between RT and practice – the "Law of Practice" – that match those observed empirically. We also show that previous methods of modelling practice in similar systems, which modulate inputs, are unlikely to successfully match observed data. The simplified versions of the model analysed are appropriate for modelling the non-stochastic parts of simple RT and two-choice RT, provide insight into the behaviour of the full version of SEEXC, and suggest a new form for the Law of Practice.

SEEXC: A model of response time in skill acquisition

In this paper we describe a model of skill acquisition (SEEXC: Self-Exciting EXpert Competition) and examine properties of simplified and mathematically tractable versions of the model. Following Usher and McClelland (1995, 2001) and Page (2000), SEEXC uses leaky competitive integration among layers of simulated neural units to model response time (RT) and accuracy. We focus on one (e.g., simple RT) and two (e.g., yes/no decision) choice cases of the model, which correspond to layers with one or two units.

Page (2000) claimed that his extension of Usher and McClelland's (1995) basic model could account for the "Power Law of Practice". In a comment on Page, Heathcote and Brown (2000a) pointed out that the "Law of Practice", the function relating a subject's mean RT to the number of practice trials (N), is much closer to an exponential function than a power function (Heathcote, Brown & Mewhort, 2000). Hence, the power function predicted by Page's model does not accord with the data.

Critically, Page (2000) assumes that practice causes an increase in the input to the model at a slower than linear rate in practice trials (N). Heathcote and Brown (2000a) claimed that a one-unit version of this input learning model could be modified to fit exponential data, but only if the input increased at an implausibly fast rate with practice. However, they also claimed that a leaky competitive integrator model could plausibly fit exponential data if practice increased each unit's self-excitation, rather than the input. SEEXC is based on the idea of practice affecting self-excitation in networks.

In this paper we prove and extend the analytic results in Heathcote and Brown (2000a). Their results apply to deterministic one-unit leaky integrators. We extend their results to deterministic two-unit leaky competitive integrators. The full SEEXC model also includes a stochastic component in the integration process (cf. Page, 2000; Usher & McClelland, 1995, 2001). We consider the deterministic case because of its mathematical tractability, and because analytic results for the deterministic case can inform the stochastic case. In particular, the results for the deterministic version provide constraint for model design and parameter selection in simulation studies of more realistic and complex networks (e.g., Brown & Heathcote, 2001; Heathcote & Brown, 2000b, 2002). Here we model only correct RT, as stochastic integration, and perhaps inputs that vary from trial to trial (cf. Ratcliff, 1978), are required in order to account for decision errors. For comparative purposes, we begin by informally examining an example of stochastic leaky competitive integration, and then provide the results for the deterministic versions.

Stochastic Leaky Integration: An Example

Figure 1 shows the evolution of the activations, $x_i(t)$, of two competing stochastic leaky integrators, where t is time. The initial values ($x_i(0)$) of both units were small and equal, and one unit received a slightly larger input (I_i) than the other ($I_1 > I_2 > 0$), although the difference was small relative to the total input ($I_1 + I_2 >> I_1 - I_2 > 0$). At first both units increased quickly in activation, responding to the large total input, so the total activation of the system ($x_1 + x_2$) quickly achieved a stable value by around t = $^1/_3$.

The further evolution of the activations was characterised by constant total activation and competition between the two units. Initially the unit indicated by the blue line, which had a slightly stronger input, gained a slight advantage, but due to the integration noise, the unit with weaker input (green line) made a come back and was briefly slightly more active around t = 1. Eventually the stronger input unit dominated and its activation increased nonlinearly, whereas the activation of weaker input unit decreased nonlinearly.

The rapidly increasing/decreasing behaviour occurs because the inhibitory competitive signal sent out by a unit is proportional to its activation. As a unit becomes more active it sends out stronger inhibition, while the less active unit sends out weaker inhibition, causing a positive feedback interaction. It is important to note that competitive signal is usually assumed to be subject to a threshold, in that negative activation values result in no competitive output (not, as might be assumed, a reversed "positive inhibition").

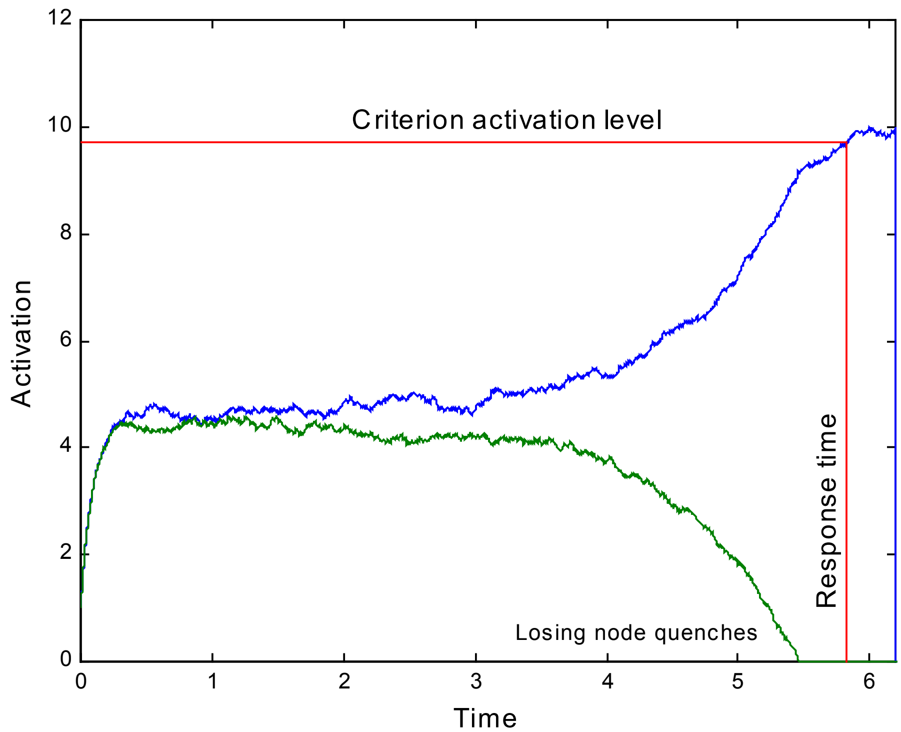

Figure 1. An example of stochastic leaky integration.

By around t = 5.4 the activation of the weaker unit was "quenched" (i.e., its activation was reduced to zero). The point at which the losing unit is quenched is called the "quenching time" (t_q). Some versions of the model allow the quenched unit to again rise above zero, but in this simulation the quenched node always remains quenched (much the same effect can be achieved by allowing the activation of the quenched node to become negative). After quenching, the dynamics of the winning node change, following the single-unit dynamics that will be examined (in a deterministic version) below.

In the example, a response was initiated shortly after quenching when the activation of the winning node crosses a criterion level. RT is equated[1] with the time at which the crossing occurs, around t = 5.8. In our development of the deterministic model, we assume that the criterion (χ) is placed so that a response occurs before quenching. This simplifies the mathematics, as there is no change from two-unit to single-unit dynamics. It is also plausible when the full stochastic version is used to model RT under speed pressure because no new information relevant to the choice accrues after quenching, making further integration redundant. In the next section we describe the SEEXC model in detail and provide analytic results for many of the features evident in Figure 1.

SEEXC

Equations 1 and 2 define the deterministic dynamics of a two-unit leaky competitive integrator with the LHS being the derivative of activation. They differ from the full SEEXC equations in that inputs are constant, rather than varying randomly from trial-to-trial, and there is no additive stochastic term. In the following we assume equal initial values $x_1(0)= x_2(0)=x_0$.

$$x_1' = I_1 - kx_1 + \varepsilon f\left(x_1\right) - \delta f\left(x_2\right) \qquad (1)$$

$$x_2' = I_2 - kx_2 + \varepsilon f\left(x_2\right) - \delta f\left(x_1\right) \qquad (2)$$

[1] Of course real choice RT data contains both decision and non-decision components. We neglect the latter here, as we are concerned with changes in decision time.

The parameters $k > 0$, representing leakage, and $\delta > 0$, representing competition, are common with Usher and McClelland's (1995) model. The output function on competitive signals, $f(x)$, is a threshold linear function, here assumed to be:

$$f(x) = \begin{cases} x & x \geq 0 \\ 0 & x < 0 \end{cases}$$

As we analyse only pre-quenching dynamics, where $x > 0$, $f(x)$ can be replaced by x.

Compared with Usher and McClelland's (1995) model, SEEXC has an extra parameter, $\varepsilon \geq 0$, representing self-excitation. It can be thought of as the weight on a recurrent self-connection, which is assumed to increase from zero to an upper bound of k with increasing practice according to the standard neural network learning law $d\varepsilon/dN = \lambda(k - \varepsilon)$, which has the solution (assuming $\varepsilon(0) = 0$):

$$\varepsilon(N) = k\left(1 - e^{-\lambda N}\right) \qquad (3)$$

As a consequence of the increase in self-excitation, the effect of practice is to increasingly offset the leakage (k) in the integration process. When leakage is large, each node's activation is controlled by recent inputs, with inputs further in the past effectively "forgotten". In contrasts, for a loss-less integrator (i.e., one in which k=ε), inputs at all times are equally weighted. For a stationary input signal, equal weighting leads to optimal information integration in the stochastic model. Brown and Heathcote (2001) described experimental evidence that is consistent with a decrease in information leakage as subjects become more expert at a task.

Sum and Difference Analysis

In order to simplify the algebra, we define sum and difference equations for the activation as (1) + (2) and (1) - (2) respectively, where $x_s = x_1+x_2$, $x_d = x_1- x_2$, $I_s = I_1+I_2$, and $I_d = I_1- I_2$:

$$x_s' = I_s - Sx_s \qquad (4)$$

$$x_d' = I_d - Dx_d \qquad (5)$$

$$\mathrm{S} = \mathrm{k} - \varepsilon + \delta \qquad (6)$$

$$\mathrm{D} = \mathrm{k} - \varepsilon - \delta \qquad (7)$$

Note that (4) and (5) hold only until quenching; that is, only while the activation of both units is strictly positive and hence $f(x) \equiv x$.

The values of S and D determine the properties of the sum and difference dynamics. Given that in SEEXC $\varepsilon \le k$, S > 0 and so the sum of activations converges to $x_s = I_S/\mathrm{S}$. Similarly, when $k > \delta + \varepsilon$, D > 0 and the difference converges to I_d/D. However, as practice increases $\varepsilon \to k$, and so D < 0 as $\delta > 0$. Hence if $k > \delta$ practice can cause a transition from converging to diverging differences, but if $\delta > k$ the difference diverges at all levels of practice. Diverging differences guarantee a response even with very small differences between inputs, and in the stochastic model a response will occur even with exactly equal inputs. Generally the sum will converge quickly relative to the difference, as illustrated in Figure 1, because typically $I_s > I_d$.

Equations (4) and (5) are examples of initial value problems of the form $dy(t)/dt = A - By(t)$, $y(0)=y_0$, which have the solutions:

$$t = \frac{-1}{B}\ln\left(\frac{A - By}{A - By_0}\right) \quad B \neq 0; \qquad t = \frac{y - y_0}{A} \quad A \neq 0, B = 0 \qquad (8)$$

From (4), (5) and (8), and using $x_s(0)=2x_0$ and $x_d(0)=0$ yields solutions for the sum (where $S > 0$ always applies) and difference equations expressed as functions of x:

$$t = \frac{-1}{S}\ln\left(\frac{I_s - Sx_s}{I_s - 2Sx_0}\right) \qquad (9)$$

$$t = \frac{-1}{D}\ln\left(\frac{I_d - Dx_d}{I_d}\right) \quad D \neq 0 \qquad t = \frac{x_d - x_0}{I_d} \quad D = 0 \tag{10}$$

Solving (9) and (10) for x in terms of t, yields:

$$x_s = \frac{I_s}{S} + \left(2x_0 - \frac{I_s}{S}\right)e^{-St} \tag{11}$$

$$x_d = \frac{I_d}{D}\left(1 - e^{-Dt}\right) \quad D \neq 0 \qquad x_d = I_d t \quad D = 0 \tag{12}$$

Since we have assumed $t < t_q$, we can use $x_1=½(x_s+x_d)$ to find an expression for activation in terms of time:

$$x_1 = \left(\frac{I_s}{2S} + \frac{I_d}{2D}\right) + \left(x_0 - \frac{I_s}{2S}\right)e^{-St} - \frac{I_d}{2D}e^{-Dt} \quad D \neq 0 \tag{13a}$$

$$x_1 = \frac{I_d t}{2} + \frac{I_s}{2S} + \left(x_0 - \frac{I_s}{2S}\right)e^{-St} \qquad D = 0 \tag{13b}$$

The equation for x_2 is the same as (13), except with I_d replaced by $-I_d$. Quenching occurs when $x_2 = 0$. Hence, t_q is the value of t for which:

$$\left(\frac{I_s}{2S} - \frac{I_d}{2D}\right) + \left(x_0 - \frac{I_s}{2S}\right)e^{-St} + \frac{I_d}{2D}e^{-Dt} = 0 \quad D \neq 0 \tag{14a}$$

$$\frac{I_d t}{2} + \frac{I_s}{2S} + \left(x_0 - \frac{I_s}{2S}\right)e^{-St} = 0 \qquad D = 0 \tag{14b}$$

A solution is available for (14) only when $x_2(t=\infty) < 0$. This condition is guaranteed when $D \leq 0$ (14a), but will only be satisfied for $D > 0$ when $I_1/I_2 > (k-\varepsilon)/\delta$, otherwise competition is not sufficiently strong to quench the losing unit, which has an asymptotic activation $x_2(t=\infty) = (I_s/S) - (I_d/D) > 0$. As noted above, all results so far are valid only for $t < t_q$.

Figure 2a shows the activations of a two-unit system with $I_s = 1$, $I_d = 0.001$, $\delta = 0.2$, $k = 0.1$ and $x_0 = 0.1$. For the red curves $N = 0$ and so $\varepsilon = 0$. They show similar behaviour to Figure 1, which was also driven by a very small input difference. The sum of activations has converged, around $t = 20$, while the difference is still small. For this system, $D = -0.1$, so the difference eventually diverges, and quenching occurs just before $t = 60$. The green curve has the same parameters except $N = 50$, which, given $\lambda = 0.2$, represents performance after extensive practice. The sum and difference increase more quickly, quenching occurs earlier, and a response would be faster, for any criterion chosen.

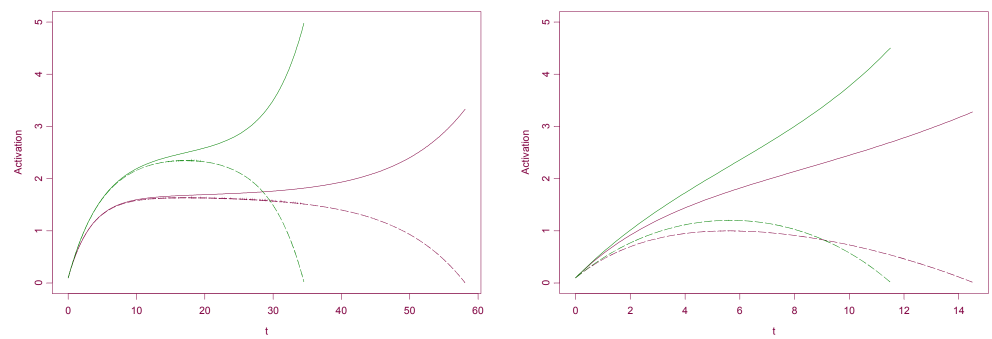

(a) (b)

Figure 2. Solid lines indicate the unit with larger input, dashed lines the unit with the smaller input, red lines indicate $N = 0$ (prior to practice) and green $N = 50$ (after extensive practice). Plots are terminated at quenching.

In Figure 2b the difference in inputs is much larger ($I_d = 0.1$), but the other parameters are the same. Activation changes more quickly and quenching occurs much earlier (note that the range of t in 2a is four times the range in 2b). The difference between each unit's activation becomes evident well before the sum has converged and the effect of practice is much smaller.

The parameters for Figure 3 are the same as for Figure 2b, except that $k = 0.2444$. The much greater leakage slows the accumulation of activation. At $N = 0$ (red lines) the difference is convergent ($D > 0$). Quenching does not occur as ε was chosen so that $I_1/I_2 = (k-\varepsilon)/\delta$. In this case the activation of the weaker input node converges to zero whereas the activation of the winning node converges to $I_1/k = 2.25$. For the blue lines $N = 1$ and $D = 0$, so integration is linear; after about $t = 10$ the non-linear effect of sum convergence is dissipated and the increase/decrease in activation is linear. Finally, the green line again represents performance after extensive practice ($N = 50$), with the difference diverging strongly. Because the initial value of k was large the effect of practice, which increases ε from zero to k, is large in this example.

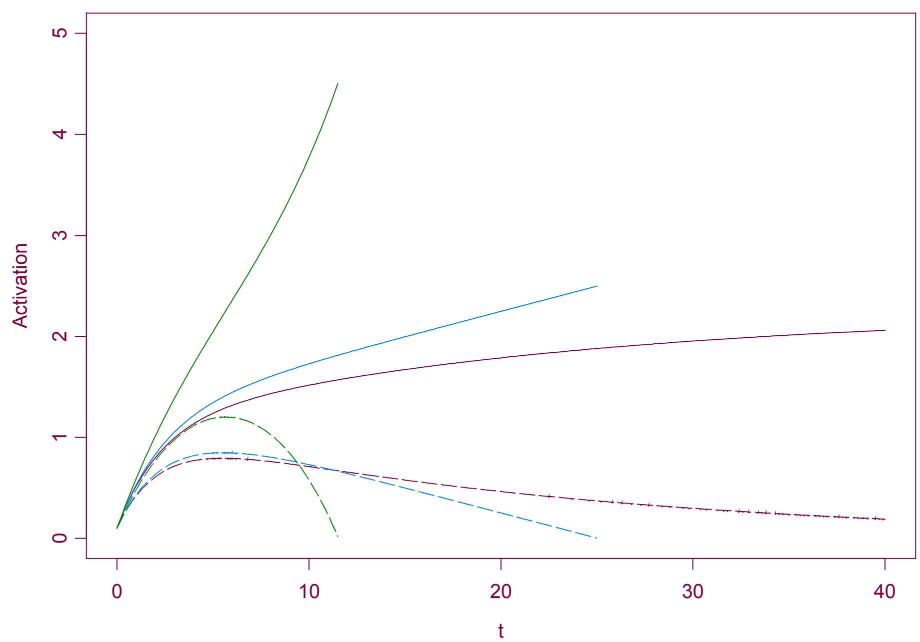

Figure 3. Solid lines indicate the unit with larger input, dashed lines the unit with the smaller input, red lines indicate $N = 0$, blue $N = 1$, and green $N = 50$. Plots are terminated at quenching, except for $N = 0$, where quenching does not occur in a finite time.

RT (assuming $I_1 > I_2$) will be the value of t such that $x_1 = \chi$, where the response criterion $\chi > x_0$ (again assuming that χ is set such that RT<t_q). This equation, and Equation (14) do not have an explicit solution for t, except in some special cases[2].

[2] For example, when $k = \varepsilon$, $S = -D$, so $x_1 = \chi$ and $x_2 = 0$ have the form $a = be^{(ct)} + de^{(-ct)}$, with solutions:

$$t = \frac{1}{c}\left(\ln \frac{a \pm \sqrt{a^2 - 4bd}}{2b}\right).$$

However, a root finding algorithm can easily determine solutions for RT and t_q numerically. The one-dimensional search entailed in finding the root is typically much quicker than obtaining RT by numerically solving (4) and (5). When $D \leq 0$, $x_1 = \chi$ always has a solution, but when D > 0, and so x_1 and x_d converge, it has a solution only when $x_1(t = \infty) = I_s/(2S) + I_d/(2D) > \chi$, which reduces to $I_1(k-\varepsilon) - I_2\delta > \frac{1}{2}\chi\left((k-\varepsilon)^2 - \delta^2\right)$.

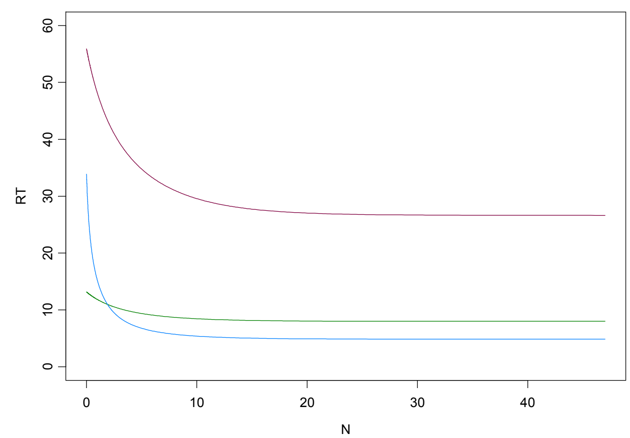

Figure 4. Practice curves corresponding to Figure 2a (red, $\chi = 3$), 2b (green, $\chi = 3$) and 3 (blue, $\chi = 2$).

Figure 4 shows the practice curves corresponding to Figure 2, with $\chi = 3$. This criterion satisfies the condition that RT < t_q, as $x_1(t_q)$ is at a minimum when $N = 0$, and when $N = 0$ $x_1(t_q) = 3.33$ for Figure 2a and $x_1(t_q) = 3.29$ for Figure 2b. Figure 4 also shows the practice curve corresponding to Figure 3, with $\chi = 2$. Note that for this case we must have $\chi < 2.25$ as x_1 does not exceed that value when N = 0. As can be seen from Figures 2 and 3, if χ is set too low there will be little difference between the two node's activations and so errors will be made in the stochastic case (cf. Figure 1). A reasonable heuristic is to set $\chi > I_s/2S$, which is half of the converged value of the sum. The lower bound is at a maximum after extensive practice, as S has its minimum value of δ, and so the condition becomes $\chi > I_s/2\delta$, which equals 2.5 for Figures 2 and

3. When the difference diverges this heuristic can always be applied, but when the difference converges, it sometimes cannot be applied, as in Figure 3. However, when $\chi = 2$ in this case the difference is large.

We now turn to a simplified case of a one-node system, where explicit solutions for RT are available, and analyse the properties of their practice functions. We then return to the two-unit system and examine some cases numerically in the light of the solutions for the one-unit system.

The Law of Practice

Heathcote et al. (2000) advocated comparison of practice curves using the relative learning rate, $RLR(N) = -RT'(N)/(RT(N) - RT(\infty))$, the rate of decrease in RT with practice divided by the remaining amount of improvement that can occur due to practice. The minus sign is included to make RLR positive for decreasing practice curves. For decreasing exponential functions RLR is a positive constant, whereas for decreasing power functions it decreases hyperbolically from infinity at $N = 0$ to zero as N increases. Heathcote et al.'s finding that an exponential function provided a better fit than the power function indicates that RLR is relatively constant. They also fit two generalizations of the power and exponential functions, the APEX function, which has an RLR that decreases from infinity at $N = 0$ to a constant greater than zero as N increases, and the generalized power function, which has a parameter for practice prior to experimental measurement. For the generalised power function RLR starts at a finite value determined by the amount of prior practice, and decreases hyperbolically to zero with further practice. The APEX function provided a better fit than the generalized power function, which Heathcote et al. interpreted as indicating that RLR remains greater than zero for large amounts of practice (i.e., as $N \rightarrow \infty$). The APEX function also fit slightly better than the exponential function, indicating that

RLR may not be constant early in practice, but may instead decrease slightly early in practice.

In order to examine the plausibility of practice functions produced by SEEXC, we now derive the RLR for some simplified versions. We begin by examining a single-unit system defined by (1) with the competition term removed. This was the case examined in Heathcote and Brown (2000a). We then examine the two-unit system numerically.

One-Unit System

The differential equation for the one-unit model, $x' = I - (k - \varepsilon)x$, can be solved yielding:

$$x(t) = \left(x_0 + \frac{I}{\varepsilon - k} \right) e^{(\varepsilon - k)t} \qquad (15)$$

This equation can be easily inverted, yielding:

$$RT = \frac{1}{\varepsilon - k} \ln\left(\frac{(\varepsilon - k)\chi + I}{(\varepsilon - k)x_0 + I} \right) \qquad (16)$$

Substituting (3) into (16) we obtain the equation for the practice curve:

$$RT(N) = \frac{1}{k} e^{\lambda N} \ln\left(\frac{I - kx_0 e^{-\lambda N}}{I - k\chi e^{-\lambda N}} \right) \qquad (17)$$

The trade-off between the exponential and logarithmic terms makes this curve difficult to study analytically. To simplify, we apply the linear Taylor approximation ln(1+A) ≅ A, and make use of the fact that:

$$\frac{I - kx_0 e^{-\lambda N}}{I - k\chi e^{-\lambda N}} \equiv 1 + \frac{ke^{-\lambda N}(\chi - x_0)}{I - k\chi e^{-\lambda N}}$$

The fraction on the right hand side is small, given reasonable choices of parameters. More importantly, it decays asymptotically to zero as N diverges to infinity. Thus, the

Taylor approximation is accurate, and becomes asymptotically perfect with increasing N. This gives the approximation:

$$RT(N) = \frac{\chi - x_0}{I - k\chi e^{-\lambda N}} \qquad (18)$$

To derive the RLR of this function, we need it's asymptotic value as N diverges:

$$\lim_{N\to\infty} RT(N) = \lim_{N\to\infty} \frac{(\chi - x_0)}{I - k\chi e^{-\lambda N}} = \frac{(\chi - x_0)}{I - k\chi \lim_{N\to\infty} e^{-\lambda N}} = \frac{\chi - x_0}{I}$$

We also need the derivative of (18) with respect to practice trials:

$$\frac{dRT}{dN} = \frac{d}{dN}\left(\frac{\chi - x_0}{I - k\chi e^{-\lambda N}}\right) = (\chi - x_0)\cdot\frac{d}{dN}\left(I - k\chi e^{-\lambda N}\right)^{-1} = \frac{(\chi - x_0)\lambda k\chi e^{-\lambda N}}{\left(I - k\chi e^{-\lambda N}\right)^2} \qquad (19)$$

Hence:

$$RLR(N) = \frac{\lambda I}{I - k\chi e^{-\lambda N}} \qquad (20)$$

The asymptotic value of this RLR function, and a value that the RLR is greater than at all times, is given by:

$$\lim_{N\to\infty} RLR = \lim_{N\to\infty} \frac{\lambda I}{I - k\chi e^{-\lambda N}} = \frac{\lambda I}{I - k\chi \lim_{N\to\infty} e^{-\lambda N}} = \lambda$$

So, the practise function specified by (18) has an RLR that is bounded, decreasing from the upper bound (at $N = 0$) to lower bound (at $N = \infty$) with practice.

$$\lambda\left(\frac{I}{I - k\chi}\right) > RLR > \lambda$$

Hence, in a one-unit system, learning on self-excitation produces a practice curve with the properties seen in data. For small values of $k\chi$ relative to I, the decrease in RLR with practice will be small, and so the practice curve will be approximately exponential.

<u>Learning on Inputs in a One-Unit System</u>

We now show that, when practice changes the input, asymptotic RLR is greater than zero only when the input increases at faster than the square of *N*. Assuming as per Page (2000) that *I* increases without bound with *N*, from (16) with $K = \varepsilon - k$ constant as a function of practice:

$$\lim_{N\to\infty} RT = \frac{1}{K}\ln\left(\frac{\lim\limits_{N\to\infty} I - Kx_0}{\lim\limits_{N\to\infty} I - K\chi}\right) = \frac{1}{K}\ln\left(\frac{I_\infty - Kx_0}{I_\infty - K\chi}\right) = 0 \qquad (21)$$

Hence, in the limit, practice makes the decision process instantaneous. The derivative of RT with respect to practice trials (N) is:

$$\frac{dRT}{dN} = \frac{x_0 - \chi}{(I - Kx_0)(I - Kx)}\frac{dI}{dN} \qquad (22)$$

Using (16), (21) and (22) the RLR is:

$$RLR(N) = \frac{K}{\ln\left(\left(\frac{I - Kx_0}{I - K\chi}\right)\left(\frac{I_\infty - K\chi}{I_\infty - Kx_0}\right)\right)}\frac{x_0 - \chi}{(I - Kx_0)(I - Kx)}\frac{dI}{dN} \qquad (23)$$

The denominator of (23) for large N is approximated by:

$$I^2 \ln\left(\frac{I - Kx_0}{I - K\chi}\right) = I^2 \ln\left(1 + K\frac{\chi - x_0}{I - K\chi}\right)$$

For large N the argument to the log function becomes approximately 1 and so the logarithm can be accurately approximated by a linear Taylor function ln(1+A) ≈ A:

$$RLR(N) = I^2 K \frac{\chi - x_0}{I - K\chi} \qquad (24)$$

This approximation becomes asymptotically perfect with increasing N, so conclusions drawn about the asymptotic behaviour of the practise curve based on the Taylor approximation are unbiased.

Equation (24) is of the order of I, so for large N the RLR of (16) is of the order:

$$\frac{1}{I}\frac{dI}{dN}$$

Thus, the asymptotic RLR goes to zero unless dI/dN increases at a faster-than-linear rate, and so asymptotically I must increase at a faster than quadratic rate. For a one-unit version of Page's (2000) model, which assumes a slower than linear increase in I with N, the asymptotic RLR will decrease to zero.

Two-Unit System

We investigated the shape of practice curves numerically for two-unit systems corresponding Figures 2 and 3. Figure 5 plots the numerically estimated RLR functions corresponding to Figure 2, using the same colours as in Figure 4. The lower curves in Figure 5 have the same parameters as the corresponding upper curves, except the learning rate is λ = 0.1 (for the upper curves it is λ = 0.2). Evidently, the RLR decreases slightly for small N and approaches λ in the limit of large N.

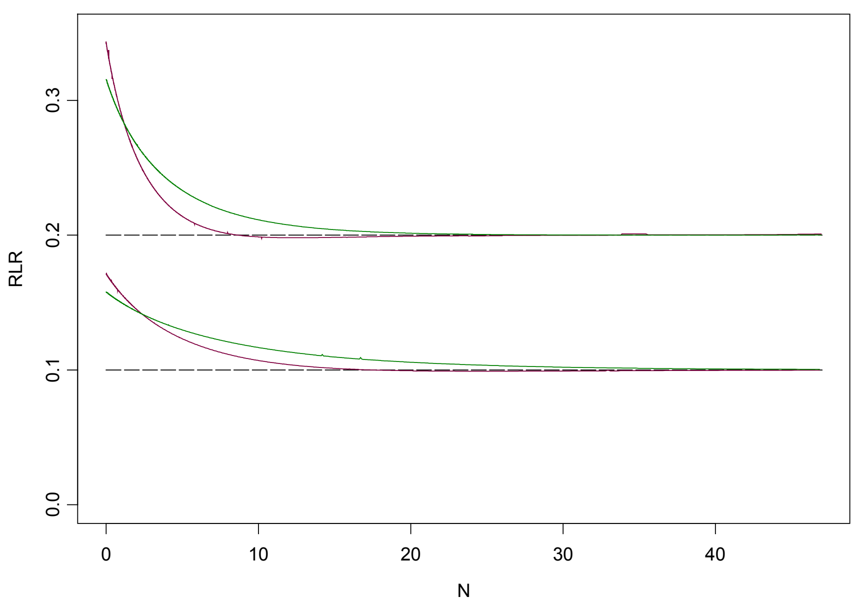

Figure 5. RLR functions corresponding to Figure 2, and Figure 2 with λ = 0.1. Values of λ are plotted as dashed lines.

Figure 6 plots the numerically estimated RLR functions corresponding to Figure 3, and Figure 3 with λ = 0.1. In this case the decrease in RLR for smaller *N* is much more dramatic, but in the limit of large *N* it still approaches λ. Hence, it appears that, at least of the cases examined here, the behaviour of the two-unit practice curves is very similar to the behaviour of the one-unit practice curves.

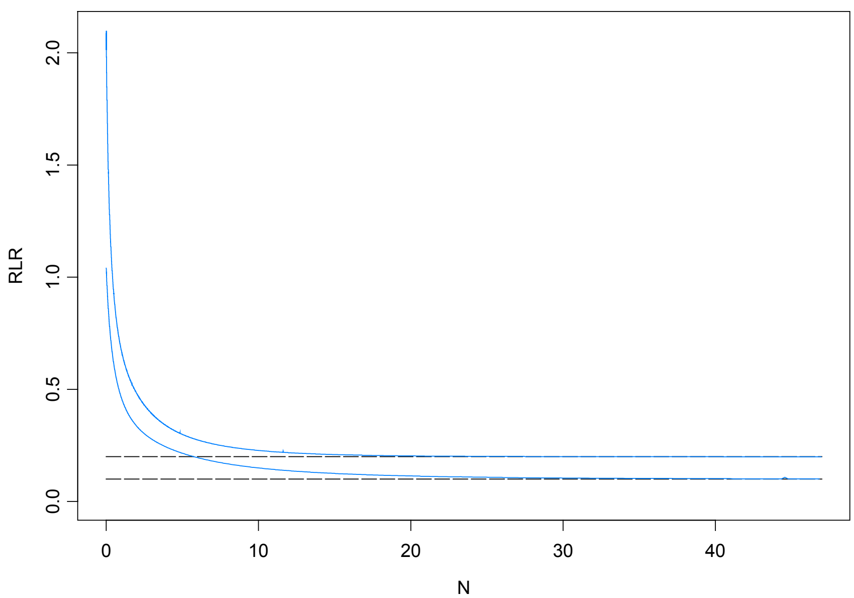

Figure 6. RLR functions corresponding to Figure 3, and Figure 3 with λ = 0.1. Values of λ are plotted as dashed lines.

Discussion

In this paper we described the SEEXC model of skill acquisition and examined the properties of deterministic versions of the model. For one-unit systems we showed analytically that learning on self-excitation could produce practice curves that are close to the form seen in Heathcote et al.'s (2000) survey of empirical practice curves. We also showed analytically that learning on inputs, as suggested by Page (2000), could only match the data when the increase in the input with practice is implausibly fast. For two-unit systems we derived the equations for the time-course of activation and identified an efficient numerical method of determining predicted RT. These results were used in a numerical investigation, which found that practice functions for two-unit systems matched one-unit systems in the properties critical for accounting for practice data. We did not investigate Page's input learning model for two-unit systems, but it seems highly likely that it will suffer similar problems to one-unit systems in matching practice data.

SEEXC is able to capture the close-to-exponential form of practice because practice changes parameters that affect activation exponentially, whereas Page's (2000) model changes the input, which affects activation linearly (e.g., Equation 13). Another key element in SEEXC's success is the neural network learning law (3). This is a standard form used so that weights remain bounded. In contrast, Page's model allows the input to increase without bound. Clearly the input signal cannot increase indefinitely, and so some modification of this assumption seems required. One method that might solve both the bounding problem, and help to match the practice function data, is to transform the input with a sigmoid non-linearity. For changes in small input values the sigmoid function increases nonlinearly, and so the rate of change of the input with practice required to produce the appropriate practice function may be achieved. However, we have not found any quantitative formulation of this type with the required properties, especially for Page's exact model, where the increase in input is slower than linear with practice.

Because our aim was to produce analytic results we examined a greatly simplified version of SEEXC. One simplification was to assume that self-excitation increased equally for all units and with the same learning rate. This seems at least a reasonable approximation where each response is performed equally often, as is commonly the case in practice paradigms. However, our results could also apply to a model where practice modulates the global level of leakage in a layer through non-local mechanisms, such as a neurotransmitter that diffuses across the entire layer. In a two-choice paradigm, these types of models make the testable prediction of equal learning rates for both responses. However, we believe that local self-excitation learning with differing rates for different units will be required in some circumstances, such as choice amongst many units that have experienced different levels of practice. Our results could also provide a "lumped" approximation to systems where each "unit" is constituted of a tightly connected network of sub-units. The activations of the sub-units are correlated because of excitatory interconnections, with the average value of the interconnection weights having a similar effect to the single self-excitation weight in the lumped approximation.

Although we think practice effects will require local leakage modulation, global modulation of leakage remains an interesting and potentially fruitful mechanism for modelling behaviour. For example, leakage could be reduced when the response is made under speed pressure, as decreased leakage causes a decision to be made more quickly. This approach contrasts with the usual mechanism of lowering the response threshold under speed pressure. Our results showed that when leakage is greater than competition the overall system is convergent, whereas for lower levels of leakage the overall system is divergent. If the system is divergent it will eventually pick a winner even when the input is not discriminative, especially in the stochastic case. Hence, high levels of leakage may be required prior to the occurrence of discriminative input to avoid the system becoming locked into a decision. Leakage could be decremented coincident with the occurrence of a discriminative input when a decision is required.

Another simplification of SEEXC was to assume deterministic integration. The stochastic and deterministic versions of the model differ in that errors, in the sense of a node with smaller input beating a node with larger input, can never occur in the deterministic version. Such errors can occur in the stochastic version because noise may allow the unit with weaker input to gain an advantage, which is then amplified by the competition, overcoming the disadvantage due to its weaker input. A second mechanism that can produce errors is random trial-to-trial variation in input strength, as assumed by Ratcliff's (1978) diffusion model, amongst others. In this case one type of "error" occurs when $I_2 > I_1$ on the error trial even though $E(I_1) > E(I_2)$, where E() indicates the expected or mean value. The distributions of the various classes of RTs (e.g., correct and error) are probabilistic mixtures of fundamental RT types, across trials. Heathcote (1998) examines the properties of such mixtures for the one-unit linear leaky integrators considered here, as well as for more complex non-linear leaky integrators. In such mixture models stochastic integration may be required to deal with cases where $I_2 \approx I_1$, which otherwise might result in no decision or extremely slow decisions.

Errors may also be produced by random trial-to-trial variation of the response criteria or the starting activations. Again for simplicity we assumed that the same criterion and starting activation was applied to each unit. If the criterion or starting activation for each unit varies randomly and at least partially independently from trial-to-trial errors can occur even in the deterministic case. Brown (2002) presents analytic and simulation results for all four sources of error, both singly and in combinations, and compares their predictions with benchmark choice RT phenomena, including the speed-accuracy tradeoff function, RT distribution and the relative speed of correct and error responses.

A limitation of our results for two-unit systems is that we only considered the case where the criterion was small enough that a response occurred before quenching. However, when a response occurs after quenching, our results for a one-unit system can be combined with our results for quenching time in the two-unit system to provide predictions. In particular, RT = t_q + RT_1, where RT_1 is the time for a one-node system to exceed the criterion when it is started at an initial value equal to the activation of the winning unit at quenching. Note that for $k > \varepsilon$ (as is always the case in SEEXC) and input I_1 to the winning unit, activation will converge to $I_1/(k-\varepsilon)$ after quenching, even when the difference diverges in the two-unit system. Consequently the converged activation of the system is bounded, for bounded inputs, and contains information about the magnitude of the input. Such information might be useful in making a confidence judgement. Often in neural network models a decision is identified with convergence rather than exceeding a criterion. Heathcote (1998) examines the use of convergence decision criteria for one-unit systems. Given that the two-unit system is convergent after quenching, convergence decision criteria may also be useful, especially in modelling confidence.

In closing we would note that the one-unit analytic results suggest a new function as a candidate for the "Law of Practice": Equation (18). For real practice data at least an additive constant (a) must be included for non-decision time, and (18) must be re-parameterised, as at least one underlying parameter is non-identifiable, resulting in:

$$RT(N) = a + \frac{b}{1 - ce^{-\lambda N}}, \qquad b = \frac{\chi - x_0}{I}, \quad c = \frac{k\chi}{I}$$

Note that b > 0, and $0 < c < 1$. As $N \rightarrow \infty$ the decision component approaches b (so RT = $a + b$), whereas when $N = 0$, the decision component is $b/(1-c)$. Under this parameterisation the bounds of the RLR can be expressed as $\lambda/(1-c) \geq RLR > \lambda$, so the decrease in RLR is entirely controlled by the *c* parameter. Further re-parameterisation produces a form with one pure location parameter, *A*, equal to RT as $N \to \infty$, and one pure scale parameter, *B*, which equals the change in RT from $N = 0$ to $N = \infty$. This form is likely to be more tractable for fitting as the location and scale parameters are both linear:

$$RT(N) = A + B\left(\frac{(1-c)e^{-\lambda N}}{1-ce^{-\lambda N}}\right), \qquad A = a + b, \qquad B = \frac{bc}{1-c}$$

Whether this new function is sufficiently similar to the actual practice functions produced by SEEXC will depend on the degree of equivalence of one and two-unit systems, the accuracy for small *N* of the linear Taylor approximation used to derive (18), and the effects of stochastic and other error-producing mechanisms. However, this "approximate SEEXC practice function" is at least of similar complexity to the four-parameter APEX and general power functions used by Heathcote et al. (2000), and instantiates the idea that RLR decreases between finite bounds.

References

Brown, S. (2002). Quantitative approaches to skill acquisition in choice tasks. Unpublished PhD Thesis, University Of Newcastle, Australia.

Brown, S. & Heathcote, A. (2001). From the O-U process to the diffusion model: Including learning and automaticity in simple decision models. Paper presented at the 10th Australasian Mathematical Psychology Conference, Newcastle University, Newcastle, Australia.

Heathcote, A. (1998). Neuromorphic models of response time, Australian Journal of Psychology, 50, 157-164.

Heathcote, A., Brown, S. (2000a). The law of practice and localist neural network models. In Page, M., Connectionist modelling in psychology: A localist Manifesto, Behavioral and Brain Sciences, 23, 479-480.

Heathcote, A., & Brown, S. (2000b) SEEXC: A self-exciting expert competitive neural network model of speed skill. Paper presented at the 33rd Annual Meeting of the Society for Mathematical Psychology, Kingston, Canada.

Heathcote, A. & Brown, S. (2002). A Sufficient Architecture for Neuromorphic Response Time Models. Paper presented at the 6th Conference of the Australasian Cognitive Science Society, Fremantle, WA, Australia.

Heathcote, A., Brown, S. & Mewhort, D.J.K. (2000) Repealing the power law: The case for an exponential law of practice. Psychonomic Bulletin and Review, 7, 185-207.

Page, M. (2000). Connectionist modelling in psychology: A localist Manifesto, Behavioral and Brain Sciences, 23, 443-512.

Ratcliff, R. (1978). A theory of memory retrieval, Psychological Review, 85, 59-108.

Usher, M. & McClelland, J. L. (1995). On the time course of perceptual choice: A model based on principles of neural computation. Technical Report PDP.CNS.95.5, Department of Psychology, Carnegie Mellon University.

Usher, M. & McClelland, J. L. (2001). The time course of perceptual choice: The leaky, competing accumulator model, Psychological Review, 108, 550-592.