

Modeling a reaction time variant of the Perruchet effect in humans

Amy McAndrew (am375@exeter.ac.uk) Fayme Yeates Frederick Verbruggen Ian P.L. McLaren

School of Psychology, College of Life and Environmental Sciences, University of Exeter, UK.

Abstract

This paper presents a reaction time (RT) experiment that follows on from the work of Perruchet, Cleeremans, and Destrebeceqz (2006), investigating the extent to which reaction times (RTs) are governed by the conscious expectancy of a particular response. In this experiment, participants were presented with a single stimulus (which we will call the conditioned stimulus; CS) followed by one of two outcomes (which we will call unconditioned stimuli; USs); to which participants had to make an appropriate instrumental response. On every trial we recorded the time taken to make this response and participants were asked to rate their expectancy that one of the USs (US1) was going to occur. We found that the expectancy rating for US1 correlated negatively with RT on US1 trials. Over successive runs of reinforcement, when participants rated US1 as less likely to occur they were slower to respond to US1 (lower ratings, higher RTs). When we calculated the expectancy for US2 as the complement of that for US1 expectancy, expectancy of US2 correlated positively with RTs. Thus, across runs of reinforcement, participants responded more quickly to US2 when considering US2 less likely (low rating, low RT). We argue that the requirement to make a conscious expectancy rating results in participants attending more to US1 occurrences than those of US2. This results in a qualitatively different relationship between conscious expectancy and automatic responses that cannot be reconciled by a single processing system account. A dual processing system explanation of learning is proposed to explain these results. In support of this position, we successfully modeled our US2 RT data using a modified version of the Augmented simple recurrent network (Yeates, Jones, Wills, McLaren, & McLaren, 2013).

Keywords: Perruchet effect; Modeling; Dual processing systems; AugSRN; Associative learning

Introduction

Recently, there has been a lively debate on the extent to which learning is governed by a single processing system or dual processing systems (e.g. McLaren, Green, & Mackintosh, 1994). A single processing system view advocates one conscious reasoning process (e.g. Lovibond, & Shanks, 2002). From this viewpoint, conditioned responding (CR) obtained in an instrumental conditioning paradigm is driven by contingency knowledge that develops during the course of conditioning between a conditioned stimulus (CS) and unconditioned stimulus (US). Within a dual processing system framework, associative automatic processes can be responsible for the CR without explicit contingency knowledge. Based on this account, an associative link forms between a representation of the CS and representation of the US. Presentation of the CS activates the link between the CS and US, which activates the US representation, which then produces a CR.

One of the most convincing sources of evidence (Mitchell, De Houwer, & Lovibond, 2009; Shanks, & St John, 1994) for dual processing systems is the Perruchet effect (Perruchet, 1985). In the reaction time (RT) version of this experiment employed by Perruchet, Cleeremans, and Destrebeceqz (2006), participants hear an auditory tone (the CS) on every trial. Half the time the CS is followed by a visual US to which participants have to make a keypress response. On the other half of the trials there is no US and participants are not required to make a response. Participants make an online expectancy rating on every trial regarding the extent to which they think the US is going to appear on that trial.

Across successive CS-US (reinforced) trials, expectancy ratings that the US will occur decrease. However, after experiencing runs of nonreinforced, CS-noUS, trials participants' ratings indicate they think it more likely that the US will occur; and thus, that a response is more likely to be required. This is consistent with the gambler's fallacy phenomenon (Burns, & Corpus, 2004). In contrast, the CR (the instrumental response to the US measured by RT) gets faster (improves) with successive reinforcement. This means consecutive CS-US trials result in shorter RTs, whereas runs involving an absence of the US result in slower responding. This pattern of responding is hard to reconcile with the gambler's fallacy, as participants become quicker to respond to the US at the same time as their expectancy of the US (and thus their expectancy that they are required to make a response) decreases. An associative account can, however, explain the change in RT with reinforcement history, as over successively reinforced trials the associative link between the CS representation and the US representation becomes stronger, leading to faster RTs. This link is extinguished and weakened by the absence of the US on the CS-noUS trials, leading to slower RTs. Thus, a dual processing systems account is required to explain both the conscious processes underlying expectancy along with the RT pattern that captures our automatic, associative learning about CS-US relationships (McLaren, Green, & Mackintosh, 1994).

The experiment presented here aims to further investigate the effects observed in a RT version of the Perruchet paradigm, and to provide support for a dual processing systems account of learning. To build on the original experiments, we presented participants with two USs in order to obtain RT data on every trial and to keep the demands of each trial consistent. We were therefore able to take a measure of CR for the two USs separately and compare these to expectancy of each US. If RT and expectancy of the US are found to follow different trends this would imply that a single processing system

explanation of learning would be unable to explain the results and that a dual processing systems account would be more appropriate. If our assumptions regarding the nature of the processes underlying RT performance are correct, we should be able to model these associatively. Therefore, to assess this claim, we used a model of human learning (the revised augmented simple recurrent network: RASRN; Yeates, Jones, Wills, McLaren, & McLaren, 2013) in an attempt to simulate the instrumental responding of participants in this experiment.

Method

Participants

64 University of Exeter students (13 men and 51 women) were recruited for course credit to participate in this experiment. Their ages ranged from 18 to 49 years, with a mean age of 21.

Design and Stimuli

The CS was visually presented to participants as a brown cylinder (11 x 7 cm) in the centre of a white screen. The words “Peanut Butter” and “Brown Sugar” were the two USs that followed the presentation of the CS. Both USs were presented to and counterbalanced across each participant as US1 and US2. Each of the USs was presented half the time after the CS, forming a partial reinforcement schedule where the occurrence of each US was equally likely.

In a typical Perruchet design, we are interested in runs of reinforced and non-reinforced trials, therefore a repeated-measures factor of run length (the number of a given trial type that occur consecutively in a row) was constructed. There were 8 levels of this factor; -4, -3, -2, -1, +1, +2, +3, and +4. When analyzing the sequence of trials given to each participant in this experiment, we can examine repetitions of the same US (D, different trials) or repetitions of the opposing US (S, same trials) as equivalents of these positive and negative runs of trials, respectively. A CR measurement is taken on the trial after the run itself, thus when considering US1 trials, a +2 trial would have involved two consecutive CS-US1 trials prior to this, whereas a -3 trial would have been preceded by a run of three CS-US2 trials (see Table 1 for an example of how runs are labeled within the sequence).

Table 1. An example of a sequence of CS-US pairings and the corresponding run lengths of these trials. These are labeled both in terms of classic Perruchet positive and negative runs; and in terms of same (S) and different (D) runs. Trial type indicates whether US1 or US2 is paired with the CS (which occurs on each trial).

Trial type	US1	US1	US2	US2	US2	US1	US2
Run		+1	-2	+1	+2	-3	-1
length		S1	D2	S1	S2	D3	D1

We aimed to compile sequences of US1 and US2 trials that involved these same (S/positive) and different (D/negative) runs from 1 to 4, following a binomial

distribution as shown in Table 2. However, the original Perruchet experiments only comprised of one CS and one US, while the current experiment involves two USs. As each run has to end in the opposite trial type (e.g. a US1 run would have to end in a US2 trial), two ‘different’ runs of length five are included in each block. These are a requirement for the sequence, are counterbalanced across the US type across blocks and excluded from the analysis; and so are not discussed further.

Table 2. The binomial distribution of run lengths.

Run length	-4	-3	-2	-1	+1	+2	+3	+4
	D4	D3	D2	D1	S1	S2	S3	S4
Number of runs	2	4	8	16	16	8	4	2

In this experiment, each participant experienced two blocks of 57 trials, which comprised of unique, randomized sequences of run lengths. These sequences were constructed using MatLab. We measured both expectancy and RT as our dependent variables and compared them across run length for both USs separately.

Procedure

A cover story was given to participants, who were told they were playing the role of a doctor seeing a number of patients with both diabetes and a nut allergy. Participants were exposed to the CS for 5 seconds on each trial and were told that this brown cylinder could represent either peanut butter or brown sugar. During this time, participants had to make a rating on a scale of 1 to 9 regarding the extent they thought this trial would be a US1 trial. For half of the participants, peanut butter was US1; for the other half, brown sugar was US1. If US1 was peanut butter, they were told that a rating of 1 would indicate: “I definitely do not think the patient will need adrenaline”; up to a rating of 9: “I definitely think the patient will need adrenaline”. Adrenaline was replaced by insulin when brown sugar was US1. Participants were told that half the time “peanut butter” would appear after the CS and on the other half of trials “brown sugar” would appear. One of these stimuli (the US) was then presented immediately after the CS. Participants were instructed to respond as quickly as possible to the stimuli to administer adrenaline to “peanut butter” and insulin to “brown sugar” with left Ctrl and left Alt keys (counterbalanced) to avoid anaphylactic shock or hyperglycemia, respectively. The US remained onscreen until a response was made, followed by a variable ITI of 2 to 5 seconds before the next trial commenced. Participants were allowed a short break between the two blocks to allow them to rest.

Results

Both RT and expectancy data were collected using MatLab and PsychToolbox (Brainard, 1997). RTs for US1 and US2 were recorded on each trial in milliseconds (ms). Any RTs over 1 second were excluded from the analyses. The mean RT for each run length for US1 and US2 can be seen in Fig. 1 top panel. In terms of expectancy, participants

were required to make ratings based on the extent they thought US1 was going to occur. Therefore, we divided the data into average expectancy for US1 on US1 trials and average expectancy for US1 on US2 trials for each participant on each run length, see Fig. 1 bottom panel.




Figure 1. The top panel displays the RT data for US1 and US2 across run length. The bottom panel displays the expectancy for US1 on US1 and US2 trials across run length.

A two-way repeated-measures analysis of variance (ANOVA) was run on the RT data using the factors US (US1 versus US2) and run length (-4, -3, -2, -1, +1, +2, +3, +4). A significant interaction between US and run length was found, $F(7,238) = 2.58$, $MSE = 0.025$, $p = .029$, as well as a significant linear trend interaction, $F(1,34) = 8.84$, $MSE = 0.085$, $p = .005$. This indicates that there is a significant difference in US1 and US2 RTs across run length. From Fig. 1 top panel, it can be seen that US1 RTs appear to increase after a run of US1 trials (i.e. RT increases when run length increases), whilst US2 RTs decrease after a run of US2 trials (i.e. RT decreases when run length increases).

One-way repeated-measures ANOVAs were then used to analyze the US1 and US2 RT data separately. There is a highly significant main effect of run length for the US2 RTs, $F(7,336) = 6.21$, $MSE = 0.07$, $p < .001$. There was also a significant linear trend decreasing from -4 to +4 across run length, $F(1,48) = 16.86$, $MSE = 0.27$, $p < .001$. With regards to US1 RTs, however, the numerically increasing linear trend from -4 to +4 was not significant.

A two-way repeated-measures ANOVA was also run on the US1 expectancy data, again with the factors US and run length. A significant interaction between US and run length was found, $F(7,371) = 3.39$, $MSE = 22.42$, $p = .017$, as well

as a significant linear trend interaction, $F(1,53) = 4.43$, $MSE = 48.92$, $p = .040$. This indicates expectancy of US1 on US1 differs significantly from expectancy of US1 on US2 trials across run length. From Fig. 1 bottom panel, it appears that expectancy for US1 on US1 trials decreases across run length whilst expectancy of US1 on US2 trials increases across run length.

One-way repeated-measures ANOVAs were then used to analyze expectancy on US1 and US2 trials separately. There is a significant main effect of expectancy of US1 on US2 trials across run length, $F(7,399) = 2.51$, $MSE = 9.78$, $p = .041$, and a significant linear trend increasing from -4 to +4, $F(1,57) = 5.38$, $MSE = 33.78$, $p = .024$. With regards to expectancy of US1 on US1 trials, a marginally significant main effect of run length was found, $F(7,392) = 2.44$, $MSE = 11.26$, $p = .051$. However, the decreasing numerical linear trend was not reliable.

Discussion

Regarding the expectancy measure (Fig. 1, bottom panel), we should make it clear from the start that both lines on the graph reflect US1 expectancy, however we have split this by whether the rating was taken on a US1 or US2 trial. Expectancy for US1 on US1 and US2 trials can be explained by the gambler's fallacy phenomenon (Burns & Corpus, 2004). Expectancy of US1 after a run of US1 trials numerically decreases, while expectancy of US1 after a run of US2 trials increases. Thus, after a run of US1 trials the participant thinks US2 is more likely to occur, so expectancy of US1 declines; but after a run of US2 trials the participant now believes it is US1s turn, so expectancy of US1 increases.

Within the RT data, participants' responses to US1 numerically increased as a function of run length. This indicates participants were faster to respond after successive CS-US2 trials, and therefore were slower after successive CS-US1 trials. We found a negative correlation between US1 expectancy and US1 RTs, $r = -.871$, $n = 8$, $p = .005$. Thus, after a run of CS-US1 trials participants made lower ratings that US1 would occur and were slower to make US1 responses. Therefore it would appear that a propositional explanation would be sufficient to explain this result, by simply claiming expectancy directly influenced RT.

Turning to US2, we propose that, logically, if a participant is expecting one US to happen then they are not expecting the other, so if a participant is expecting a US1 trial to occur then that implies they are not expecting a US2 trial. This would suggest expectancy of the two USs is complementary such that, if expectancy of US1 is the highest possible rating (9), then expectancy of US2 should be the lowest possible rating (1). We can assume that these sum ($9+1=10$) and thus calculate expectancy for US2 as equal to 10 minus US1 expectancy. Based on this assumption, we can predict participants' expectancy of US2 on US2 trials as being the complement of their expectancy of US1 on US2 trials, see Fig. 2. If this supposition is true, then we have shown expectancy of US2 on US2 trials decreases as a function of run length. Therefore, higher ratings of US2 are made if participants have experienced a

run of US1 trials, and vice versa. This pattern of responding can be attributed to the propositional, gambler's fallacy phenomenon discussed previously.




Figure 2. This graph displays expectancy of US1 on US2 trials and the hypothetical expectancy of US2 on US2 trials.

In order to verify if our inference regarding expectancy of US2 on US2 trials was correct, 32 of our participants carried out a further two experimental blocks to those described in the earlier method section. In these blocks, two (identical) cylinders were presented (successively) and participants had to make an expectancy rating to each. One cylinder required the participants to make a “peanut butter” rating, the other a “brown sugar” rating. Participants then had to make the appropriate RT response as in the previous blocks. Comparing participants expectancy of US1 on US1 trials and their expectancy of US2 on US1 trials, there was a highly significant negative correlation, $r = -.969$, $n = 8$, $p < .001$. This shows that on US1 trials, if participants were for example, expecting a US1 trial they were not expecting a US2, and vice versa. Additionally, comparing expectancy of US1 on US2 trials and expectancy of US2 on US2 trials, there was also a highly significant negative correlation, $r = -.944$, $n = 8$, $p < .001$. This also shows that on US2 trials, if participants were expecting a US1 they were not expecting a US2. Therefore, our earlier assumption receives considerable empirical support from this check.

Given that expectancy of US2 on US2 trials decreases as a function of run length, interestingly we found that US2 RTs also decreased as a function of run length (see Fig. 1). Participants were faster to respond to US2 on a run of CS-US2 trials, even though their expectancy that US2 would occur had decreased. We have therefore demonstrated a positive correlation between expectancy of US2 on US2 trials and US2 RTs, $r = .833$, $n = 8$, $p = .010$. For example, after a run of CS-US2 trials, participants rate that a US1 trial is more likely (and therefore a US2 is less likely), yet are faster to respond to US2. It is consequently hard to reconcile this expectancy with the RT result if we take the position that a single propositional explanation could explain our data. We would argue that associative, link-based processes are required to explain the RTs for US2. One version of this would be that when a person experiences the CS followed by US2, a link is set up between the two representations of these stimuli. After a run of CS-US2 trials this would strengthen the link between these stimuli, resulting in a stronger CR (i.e. a faster key press response) to US2.

However, after a run of CS-US1 trials, the link between CS and US1 strengthens, but the link between the CS and US2 weakens (extinction). Hence, the more consecutive CS-US1 trials there are, the weaker the CR to US2 (i.e. the slower the RT). The results for US2 are in agreement with previous Perruchet RT experiments, in which a single propositional process cannot explain both the expectancy and RT data.

In one experiment we have shown two different results, one where expectancy and RT appear positively correlated, and another where they are negatively correlated. We have, as a consequence, proposed a dual processing systems explanation of the US2 result. We would now like to pursue this further, by speculating how associative and propositional processes could produce both results. We hypothesize that the difference between the two effects (for US1 and US2) lies in where participants' attention is focused. As participants are directed to focus on one US (US1), to which they are making expectancy ratings, this effectively manipulates the expression of both propositional and associative processing systems for that US. We assume that because participants are attending to US1, they spend less time thinking about US2 and this would suggest conscious reasoning processes are more focused on the processing of US1 than US2. If US2 is not being consciously processed (to the same extent) then changes in US2 performance in the experiment might be driven by an alternative processing system. By reducing attention to US2, we believe we have created an environment conducive to associative learning. In contrast, a large amount of cognitive resource is being directed to processing US1, and perhaps this has led conscious processes to play a larger role in RT performance for this outcome, and inhibited the expression of associative processes in this case.

Modeling

To explore how associative processes might be driving instrumental responses to US2, we chose to simulate this experiment using an established model of associative learning. We chose the augmented SRN (Cleeremans, & McClelland, 1991; as revised by Yeates, Jones, Wills, McLaren, & McLaren, 2013), which is particularly well-suited to this task as the simple recurrent network (SRN; Elman, 1990) was devised to account for learning that is observed across sequences of trials. Our aim was to ascertain the extent to which learning is driven by the development of associative strength between the CS and US2, or whether the sequential structure of the experiment (runs of US1 and US2) is what drives this result.

The model (see Fig. 3) is a connectionist network that feeds input activation to a hidden layer, which in turn feeds activation forward into an output layer, each employing the logistic activation function (Rumelhart, Hinton, & Williams, 1985). The activation of the hidden layer is copied back into a set of context units on each trial, which are then fed into the hidden layer as input on the next trial. This recurrent loop provides the model with a memory of the hidden layer's representation of the last trial. Learning occurs through back-propagation of error correction, comparing output activation to expected responses. Connection weight

changes to represent both short- and long-term learning are calculated using fast and slow learning rates, respectively. Fast weights have a higher learning rate but decay more rapidly, and were introduced to the model by Cleeremans and McClelland (1991) to account for the short term priming effects evident in their data. The slow weights reflect more permanent learning that takes a longer time to develop due to the lower learning rate.




Figure 3. Architecture of the revised version of the Augmented SRN by Yeates et al. (2013)

The model in this simulation involved two output units to represent Ctrl and Alt keypress responses to US1 and US2. As well as the context units (copy of the previous trial hidden unit activation) there were five additional input units. These followed revisions to the SRN by Yeates et al. (2013, see for further discussion) and included both a representation of the previous response made (two units, one for US1 and one for US2) as well as a representation of the on-screen stimuli on the current trial (one CS unit and two US units, one to represent each of US1 and US2). The free parameters of the model were: 20 hidden units with the learning rates set at 0.4 and 0.533 for slow and fast learning rates, respectively (based on Jones, & McLaren, 2009).

The model was run 64 times with random initial weights of between -0.5 and 0.5 to give the same n of networks as participants in the experiment. Each of these simulations used binary input activations representing the exact occurrence of the CS and USs taken from the unique sequences that each of the 64 participants were given. Mean square error (MSE) was calculated as an index of responding to the US on each trial from the squared difference between output activations and the expected activations for the two possible responses (0.1 and 0.9 for incorrect and correct response, respectively). Trials were analyzed according to run length and US, like the variables of interest used in the behavioral experiment.

We analyzed the MSE for each US using one-way repeated measure ANOVAs and thus examined the modeling data in the same fashion as the behavioral data. There was a main effect of run length in both US1, $F(7,406) = 1339.80$, $MSE = 0.67$, $p < .001$, and US2, $F(7,441) = 1546.46$, $MSE = 0.67$, $p < .001$. Thus, for both US1 and US2 MSE differed according to run lengths. Furthermore, we found that there was a highly significant linear contrast on run length for both USs, $F(1,58) = 2633.43$, $MSE = 4.44$, p

$< .001$, and $F(1,63) = 2908.722$, $MSE = 4.14$, $p < .001$, for US1 and US2 respectively. This is seen quite clearly in Fig. 4, which shows a decreasing linear trend for both USs (which do not differ significantly) across run length. It can also be seen from the graph the two functions of MSE lie almost entirely on top of one another. Thus, responding to both of these USs is extremely similar, both demonstrating a reduction in error as run length increases.




Figure 4. Graph of the mean square error (MSE) of the model

When comparing the modeling data to the human data we are using MSE as an approximation to RTs, as this is what we consider to capture the automatic, associative relationship between CS and US. We can see that human RT responding to US2 has the same, decreasing function across increasing reinforcement as is produced by the AugSRN. This is supported by a significant positive correlation between run length on RT and MSE results for US2, $r = .895$, $n = 8$, $p = .003$. Clearly then, the Augmented SRN is a good model of human performance on US2 in our experiment, but a poor one for US1.

Further investigation, however, reveals that the basis for performance may not be the conventional associative explanation offered for the Perruchet effect. There is no doubt that transient fluctuations in the strength of CS-US associations could explain the results observed for US2. But, the Augmented SRN can also learn about the sequence of events that take place, rather than just in terms of CS-US associations; and with the parameters given in Yeates et al., (2013) it could be that the pattern shown in Fig. 4 is based on this type of learning, rather than CS-US learning. This can be investigated by running the same simulation, but with the CS unit permanently set to zero so that no change in CS-US associations is possible. When we did this, the same function emerged, see Fig. 5. Thus, we would appear to have evidence suggesting that transient changes in CS-US associations might not be the basis of the function shown in Fig. 4. This result is reminiscent of that reported by Mitchell, Wardle, Lovibond, Weidemann and Chang (2010), who were able to get a Perruchet type effect in an RT experiment without any CSs. We have essentially the same result in our simulation, using a model that is well known for its ability to generate sequential effects.

But if sequential effects are the correct explanation of our modeling result, the removal of all the input units from the model (leaving only the hidden and output layers) should

abolish this effect, as there would be nothing left in the model that could produce sequential effects (no input or copy-back from the hidden layer). However, when we did this, we found the same decreasing function in MSE as seen before in our previous simulations (see Fig. 5). This demonstrates that sequential effects are not necessarily driving our result, but rather that the associative fluctuations between the hidden and output units are.




Figure 5. Graph of the MSE for the further modelling

At the beginning of these last simulations the hidden units have activation values of 0.5 (corresponding to zero input). Therefore, after a reinforced trial the link between any hidden unit and the output unit will be strengthened. Consequently, if another reinforced trial follows the previous one this link is again strengthened leading the model to produce a smaller MSE. In contrast, a nonreinforced trial weakens this link, and the MSE increases. Therefore, there is an associative explanation for the Perruchet effect that emerges from this model, just not the classic explanation as it is usually cited. It is worth emphasizing that it is an associative explanation that applies here, and not one based on conscious, cognitive expectancy of the US. The pattern seen for US1 in our empirical data follows that generated by the expectancy ratings given by our participants and is quite different from both the pattern seen for US2 and the pattern generated by our model simulating an explanation in terms of CS-US associations, sequential effects, or hidden to output layer connections. The correlation between human RTs and modeling data for US1 is negative and non-significant across run length, $r = -.562$, $n = 8$, $p = .148$. Thus, an associative explanation will not fit these data, and a more cognitive model is required.

General discussion

This paper presents behavioral and modeling data based on a new RT variant of the classic RT Perruchet paradigm. In our behavioral experiment we produced a Perruchet-type effect whereby expectancy of US2 decreased as a function of run length while RT responses to US2 decreased. We have rejected a single processing system explanation of learning in favor of a dual processing systems argument to explain this result. The propositional, gambler's fallacy heuristic (Burns & Corpus, 2004) explains why expectancy of US2 decreased as the run of CS-US2 trials increased, as participants are deciding that it is less likely another US2 trial will happen if they have experienced a run of US2

trials. However, within the RT data, after a run of CS-US2 trials participants are faster to respond to US2 despite low expectancy that US2 will occur. This seems paradoxical when considered from a single systems view, but an associative account can explain the RT result, in terms of fluctuating hidden-output unit associations, sequential effects or CS-US associations. Our feeling is that it would be possible to parameterise the Augmented SRN to produce the US2 pattern of results on the basis of any of these potential mechanisms, though it would appear that in our current simulations the effect is mainly carried by fluctuating hidden-output associations. Note, however, that in Fig. 5 the pattern is more pronounced when the input to the model is enabled (suggesting that sequential effects can contribute), and we have run other simulations that show that the presence or absence of a CS representation can also strengthen or weaken this effect indicating that CS-US associations can also be effective in this model. More research will be needed to determine which of these mechanisms is the correct explanation for our data.

In conclusion, the evidence for dissociable propositional and associative processes provided by Perruchet type RT experiments is perhaps stronger than we thought. Explaining these effects with reference to a single propositional system, however, is likely to prove a difficult challenge for theorists of that persuasion.

ACKNOWLEDGEMENTS

This research was funded by an ESRC studentship to Amy McAndrew and an ESRC grant to Ian McLaren.

REFERENCES

- Brainard, D.H. (1997). The Psychophysics Toolbox. *Spatial Vision*, 10, 433-6.
- Burns, B.D., & Corpus, B. (2004). Randomness and inductions from streaks: "Gambler's fallacy" vs. "Hot hand". *Psych. Bulletin and Review*, 11, 179-184.
- Cleeremans, A., & McClelland, J.L. (1991). Learning the structure of event sequences. *Journal of Experimental Psychology: General*, 120, 235-253.
- Elman, J.L. (1990). Finding structure in time. *Cognitive Science*, 14, 179-211.
- Jones, F.W., & McLaren, I.P.L. (2009). Human sequence learning under incidental and intentional conditions. *Journal of Experimental Psychology: Animal Behavior Processes*, 35, 538-553.
- Lovibond P.F., & Shanks, D.R. (2002). The role of awareness in Pavlovian conditioning: Empirical evidence and theoretical implications. *Journal of Experimental Psychology: Animal Behavior Processes*, 28, 3-26.
- McLaren, I.P.L., Green, R.E.A., & Mackintosh, N.J. (1994). Animal learning and the implicit/explicit distinction. In N.C. Ellis (Ed.) *Implicit and explicit learning of languages* (pp. 313-332). New York, NY: Academic Press.
- Mitchell, C.J., De Houwer, J., & Lovibond, P.F. (2009). The propositional nature of human associative learning. *Behav. and Brain Sciences*, 32, 183-246.
- Mitchell, C.J., Wardle, S.G., Lovibond, P.F., Weidemann, G., & Chang, B.P. (2010). Do reaction times in the Perruchet effect reflect variations in the strength of an associative link? *Journal of Experimental Psychology: Learning Memory and Cognition*, 36, 567-72.
- Perruchet, P. (1985). A pitfall for the expectancy theory of human eyelid conditioning. *Pavlovian Journal of Biological Sciences*, 20, 163-170.
- Perruchet, P., Cleeremans, A., & Destrebecqz, A. (2006). Dissociating the effects of autonomic activation and explicit expectancy on reaction times in a simple associative learning task. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 32, 955-965.
- Rumelhart, D.E., Hinton, G.E., & Williams, R.J. (1985). Learning internal representation by error propagation. In D. E. Rumelhart & J. L. McClelland (Eds.), *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*. (Vol. 1: Foundations.). Cambridge, MA: Bradford Book/MIT Press.
- Shanks, D.R., & St. John, M.F. (1994). Characteristics of dissociable human learning systems. *Behavioural and Brain Sciences*, 17, 367-447.
- Yeates, F., Jones, F.W., Wills, A.J., McLaren, R.P., & McLaren, I.P.L. (2013). Human sequence learning under incidental conditions. *Journal of Experimental Psychology: Animal Behavior Processes*, 39(2), 166-173.