## Brief Communication

# Context specificity of both acquisition and extinction of a Pavlovian conditioned response

Sarah Starosta,[1,4] Metin Uengoer,[2] Isabelle Bartetzko,[1] Sara Lucke,[2] Onur Güntürkün,[1] and Maik C. Stüttgen[3]

[1]Faculty of Psychology, Department of Biopsychology, Ruhr University Bochum, 44780 Bochum, Germany; [2]Faculty of Psychology, Section for Experimental and Biological Psychology, Philipps-University Marburg, 35037 Marburg, Germany; [3]Institute of Pathophysiology & Focus Program Translational Neurosciences, University Medical Center of the Johannes Gutenberg University Mainz, 55128 Mainz, Germany

It is widely held that the extinction of a conditioned response is more context specific than its initial acquisition. One proposed explanation is that context serves to disambiguate the meaning of a stimulus. Using a procedure that equated the learning histories of the contexts, we show that the memory of an appetitive Pavlovian association can be highly context specific despite being unambiguous. This result is inconsistent with predictions of the Rescorla–Wagner model of learning but in line with configural accounts of contextual control of behavior. We propose an explanatory model in which context serves to modulate the gain of associative strength and which expands upon the configural idea of unitary representations of context and conditioned stimuli.

[Supplemental material is available for this article.]

Many decades of research indicate that extinction of a conditioned response does not erase the excitatory memory of acquisition (for review, see Todd et al. 2014). One line of evidence for this claim comes from studies showing that extinction is disrupted by contextual manipulations. For instance, when acquisition and extinction are conducted in two different contexts, extinguished responses reappear when tested in the context of initial acquisition (ABA renewal; Bouton and King 1983) or in a novel context (ABC renewal; Rescorla 2008). Response recovery was also observed when acquisition and extinction took place in the same context and testing in a different one (AAB renewal; Rescorla 2008). These various forms of renewal indicate that extinction performance is more context specific than initial acquisition (Bouton et al. 2006; Rosas et al. 2013).

Learning theories differ in their assumptions about the mechanisms underlying response recovery following extinction. According to Bouton (2004), extinction equips the conditioned stimulus (CS) with a second association that signals the absence of the unconditioned stimulus (US) counteracting the excitatory CS–US association established during acquisition. Retrieval of the second-learned association requires the presence of the context of extinction, while activation of the first-learned association proceeds independently of context. Bouton's retrieval model is able to account for ABA, ABC, and AAB renewal but is challenged by demonstrations of context specificity of simple acquisition learning (Hall and Honey 1990; Üngör and Lachnit 2006). A second account is provided by Rescorla and Wagner (1972). The theory assumes that the contexts of acquisition and extinction acquire direct excitatory and inhibitory associations with the US, respectively. The inhibitory impact of the extinction context is predicted to "protect" the CS from a complete loss of its excitatory associative strength. A third account is offered by Pearce's

(1994) configural theory assuming that a specific combination of context and CS results in one unitary representation that develops an association to the US. The theories proposed by Rescorla and Wagner (1972) and Pearce (1994) are able to account for ABA and ABC renewal and for context specificity of acquisition. However, the theories cannot be easily applied to the phenomenon of AAB renewal.

The aim of the present experiment was to investigate the mechanisms underlying the context specificity of appetitive sign-tracking behavior by differentiating between the abovementioned theories. We used pigeons as experimental animals due to their inherent ability to work on several (visual) stimuli in parallel and their sensitivity to reward contingencies (Stüttgen et al. 2011, 2013; Starosta et al. 2013, 2014; Güntürkün et al. 2014). To differentiate the Rescorla–Wagner model from the theories proposed by Bouton (2004) and Pearce (1994), we equated the learning histories of the contexts according to a procedure proposed by Rescorla (2008; see also Lengersdorf et al. 2014, 2015). Under these conditions, the Rescorla–Wagner theory predicts perfect generalization of extinction across contexts, while the other two accounts anticipate renewal. To differentiate between Bouton's (2004) and Pearce's (1994) theories, we investigated context specificity of both acquisition and extinction. According to Pearce (1994), performance in both cases should be context specific, while Bouton (2004) anticipates context specificity only related to extinction.

Animals ($n = 18$) were trained in two conditioning chambers; the absolute number of conditioned responses (pecks directed onto the visual conditioned stimuli) served as the main dependent variable (see Fig. 1 for a sketch of the paradigm).

[4]Present address: Cold Spring Harbor Laboratory, 1 Bungtown Road, Cold Spring Harbor, NY 11724, USA.
Corresponding author: sstarost@cshl.edu

Phase 1

Phase 2

Phase 3

ABA          ABB          BA          BB
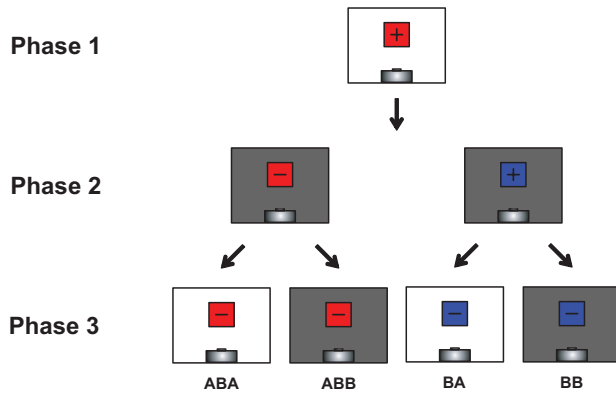
**Figure 1.** Depiction of the within-subject ABA renewal design. Single pictures illustrate physical contexts in the different phases of the experiment. The actual wallpapers used in the conditioning chambers are shown in a previous paper from our group (Lengersdorf et al. 2014). The red squares with plus signs indicate the conditioned stimulus in the first phase where stimulus presentation was followed by food. In the second phase, this stimulus was presented in a different context and was no longer followed by food (minus signs). In addition, a novel stimulus was presented and followed by food presentation (blue square with a plus). Finally, all stimuli were tested in extinction in both contexts (phase 3). Not shown are the target stimulus (present and reinforced in all sessions) and the nontarget stimulus (present and nonreinforced in all sessions). Note that only one path of the within-subject design is shown here (e.g., CS1-A and CS2-A). To display both conditions, the current drawing needed to be duplicated with swapped background colors (e.g., Lengersdorf et al. 2014). Sequence of contexts tested was balanced across subjects.

During Phase 1, animals acquired responding to CS1-A in context A and to CS1-B in context B. Then, during Phase 2, responding to each CS was extinguished in the alternate context (CS1-A in context B and CS1-B in context A). In this second phase, a novel stimulus was introduced in each context and was consistently followed by food (CS2-A in context A and CS2-B in context B). Finally, we tested responding to all stimuli in both contexts in extinction (Phase 3). This within-subject procedure ensured that animals were exposed to both contexts equally often and received the same numbers of stimulus presentations and reinforcements in both contexts (see Supplementary Methods for details). Where appropriate, we collapsed responding to CS1-A and CS1-B and simply refer to CS1, and collapsed responding to CS2-A and CS2-B and refer to CS2. In case of CS1, the respective acquisition and extinction contexts are referred to as context A and context B, respectively. In case of CS2, the respective acquisition and test contexts are referred to as context B and context A, respectively.

To determine whether pooling of these two cases (CS1-A and CS1-B) is reasonable and to confirm successful acquisition of responding to the CS1s, we compared responding to the CS1s with response levels to a nonreinforced control stimulus. Differential responding to two control stimuli (one consistently reinforced [target, T] and one nonreinforced [nontarget, NT]) was established in a pretraining phase (Table 1; see Supplementary Methods). In Figure 2A, animals showed differential responding to reinforced and nonreinforced stimuli in the last session of Phase 1, independent of the context in which the stimuli were presented. This was confirmed by a repeated-measures ANOVA (rmANOVA) with factors stimulus (NT, CS1) and context (A, B) which yielded a main effect of stimulus ($F_{(1,17)} = 60.55$, $P < 0.001$, $\eta^2 = 0.60$), but no effect of context ($F_{(1,17)} = 0.05$, $P = 0.82$, $\eta^2 < 0.01$) and no interaction ($F_{(1,1)} = 0.02$, $P = 0.87$, $\eta^2 < 0.01$). A paired $t$-test revealed no difference in responding between CS1-A and CS1-B ($t_{17} = 0.20$, $P = 0.84$, $g = 0.03$). Therefore, we pooled responses to CS1-A and CS1-B in all further analyses and will consistently use the term "CS1" to refer to both CS1-A and CS1-B.

In the second experimental phase (Fig. 2B), the two CS1s were presented in a different context than during acquisition (CS1-A in context B and CS1-B in context A) and were no longer followed by food. In addition, another consistently reinforced stimulus was introduced in each context (CS2-A and CS2-B, summarized as CS2). Thus, acquisition of conditioned responding to CS2 occurred in the same session as extinction of responding to CS1. As confirmed by an rmANOVA with factors stimulus (CS1, CS2) and block, we observed a difference between responding to CS1 and responding to CS2 over time (main effect of stimulus ($F_{(1,17)} = 28.43$, $P < 0.001$, $\eta^2 = 0.23$), reflecting more responses to CS2 than to CS1; main effect of block ($F_{(9,17)} = 9.36$, $P < 0.001$, $\eta^2 = 0.05$); stimulus–block interaction ($F_{(9,1)} = 26.67$, $P < 0.001$, $\eta^2 = 0.1$)). Separate one-way rmANOVAs with factor block yielded significant effects as well (CS1: $F_{(9,17)} = 22.35$, $P < 0.001$, $\eta^2 = 0.46$; CS2: $F_{(9,17)} = 3.47$, $P < 0.001$, $\eta^2 = 0.02$). Conditioned responding to CS2 was acquired rapidly, as indicated by the shallow increase over the first four blocks of trials (Fig. 2B, inset; also see Supplementary Results).

The main finding of the study is illustrated in Figure 2C,D: When subjects received nonreinforced presentations of CS1 and CS2 in both contexts A and B, responding to CS1 was stronger in context A while responding to CS2 was stronger in context B, as indicated by a significant stimulus–context interaction ($F_{(1,1)} = 37.68$, $P < 0.001$, $\eta^2 = 0.20$) of an rmANOVA (stimulus by context). Context-specific responding was also confirmed by post hoc $t$-tests between response counts in A and B for both stimuli (CS1: $t_{17} = 5.42$, $P < 0.001$, $g = 1.58$; CS2: $t_{17} = 5.22$, $P < 0.001$, $g = 1.19$).

Thus, we observed both ABA renewal and context specificity of simple acquisition in the present experiment. Either of these findings is inconsistent with the learning theory of Rescorla and Wagner (1972) as, in the present procedure, the contexts were equated for their associative learning histories. The results are also inconsistent with Bouton's (2004) theory. While the model does predict response recovery following extinction as observed here, the observation of context-specific acquisition performance contradicts the hypothesis that contextual cues are only integrated into the memory of ambiguous stimuli.

Our results are consistent with the configural theory proposed by Pearce (1994). According to his model, acquisition and extinction conducted in different contexts establish two unitary representations, each encoding the specific configuration of context and CS. The response-eliciting property of a configuration is determined by its direct association to the US and by generalized associative strengths of other configurations based on similarity.

One way to reconcile Bouton's (2004) retrieval model with the present findings would be to adopt an extension of the theory suggested by Rosas et al. (2006). These authors assumed that increased processing of context due to extinction results in all the information being processed in a way that makes it context specific, regardless of whether acquisition or extinction learning is involved. Future research may differentiate the extended model by Rosas et al. (2006) from Pearce's (1994) theory by investigating

**Table 1.** Overview of the training and testing phases

| Phase | Context | Target | Nontarget | CS1-A or CS1-B | CS2-A or CS2-B | Trials per session |
|---|---|---|---|---|---|---|
| Pretraining I | A | 60 (+) | | | | 60 |
| | B | 60 (+) | | | | 60 |
| Pretraining II | A | 32 (+) | 30 (−) | | | 62 |
| | B | 32 (+) | 30 (−) | | | 62 |
| Phase 1 | A | 26 (+) | 24 (−) | 24 (+) | | 74 |
| *5 sessions each* | B | 26 (+) | 24 (−) | 24 (+) | | 74 |
| Phase 2 | A | 17 (+) | 15 (−) | 30 (−) | 30 (+) | 92 |
| *2 sessions each* | B | 17 (+) | 15 (−) | 30 (−) | 30 (+) | 92 |
| Phase 3 | A | 44 (+) | 12 (−) | 12x CS1a (−) and 12x CS1b (−) | 12x CS2a (−) and 12x CS2b (−) | 104 |
| *1 session each* | B | 44 (+) | 12 (−) | 12x CS1a (−) and 12x CS1b (−) | 12x CS2a (−) and 12x CS2b (−) | 104 |

Number of trials per session in each context for every phase of the experiment. Plus signs indicate reinforced stimulus presentations, minus signs nonreinforced presentations.

context specificity of acquisition learning in the absence of any concurrent form of extinction treatment (for other possibilities, see Rosas and Callejas-Aguilera 2006; Lucke et al. 2013).

In previous studies demonstrating context-specific acquisition performance, the context-shift effect was often ascribable to differences in the associative values of the contexts (Lovibond et al. 1984; Hall and Honey 1990). Evidence for context specificity of acquisition with equated learning histories of the contexts was provided by Rescorla (2008), but his conclusion was based on a cross-experimental comparison. The present findings demonstrate the reliability of Rescorla's conclusion by providing strong evidence that contextual control of acquisition does not require differences in the contexts' learning histories.

It is widely assumed that context specificity of responding results from generalization decrement of performance after a context switch. However, previous studies reporting a context-switch effect for acquisition performance tested responding in extinction and averaged across several trials (Rescorla 2008; Bouton and Todd 2014) which may shadow perfect generalization in the first trials. We were interested in the development of context-specific responding and analyzed responding on a trial-by-trial basis. Figure 2D plots the average number of responses to CS1 and CS2 in contexts A and B for each trial. For CS1, a clear effect of context can be observed already in the first trial (rmANOVA, main effect of context ($F_{(1,17)} = 27.20$, $P < 0.001$, $\eta^2 = 0.14$); main effect of trial ($F_{(11,17)} = 9.77$, $P < 0.001$, $\eta^2 = 0.14$); context–trial interaction

($F_{(1,11)} = 2.78$, $P = 0.002$, $\eta^2 = 0.02$). Post hoc *t*-tests revealed significant differences between the contexts within the first six but no other trials (all $P$'s < 0.05, $0.95 <$ all $g$'s < 1.66). For CS2, both context and trial effects were observed as well. The rmANOVA yielded a main effect of context ($F_{(1,17)} = 10.00$, $P < 0.001$, $\eta^2 = 0.15$) as well as trial ($F_{(11,17)} = 29.33$, $P < 0.001$, $\eta^2 = 0.14$), as well as a context–trial interaction ($F_{(1,11)} = 7.53$, $P < 0.001$, $\eta^2 = 0.06$), reflecting more responding in context B and a decrease over time. However, different from CS1, responding did not differ between contexts in the very first trial ($t_{17} = 0.37$, $P = 0.72$, $g = 0.08$) but from the second trial on until the end of the session (all $P$'s < 0.05, $0.52 <$ all $g$'s < 1.52). Thus, our detailed analysis demonstrates full generalization on the first trial for CS2 and earlier extinction in a novel context rather than immediate generalization decrement after the context switch. This is supported by similar response levels in the last trial of acquisition for CS2 and the first trial of retrieval in the novel context (context switch from A to B: $t_{17} = 0.12$; $P = 0.9$; $g = 0.29$; switch from B to A: $t_{17} = 0.34$; $P = 0.74$; $g = 0.37$). Critically, we also observed perfect generalization for CS1 following the context change from acquisition to extinction, as indicated by similar levels of responding in the last trial of the last acquisition session and the first trial in extinction (switch from A to B: $t_{17} = 1.36$; $P = 0.19$; $g = 0.29$; switch from B to A: $t_{17} = 1.40$; $P = 0.17$; $g = 0.37$). Importantly, both imperfect generalization and earlier extinction imply contextual control of responses as it was pointed out in
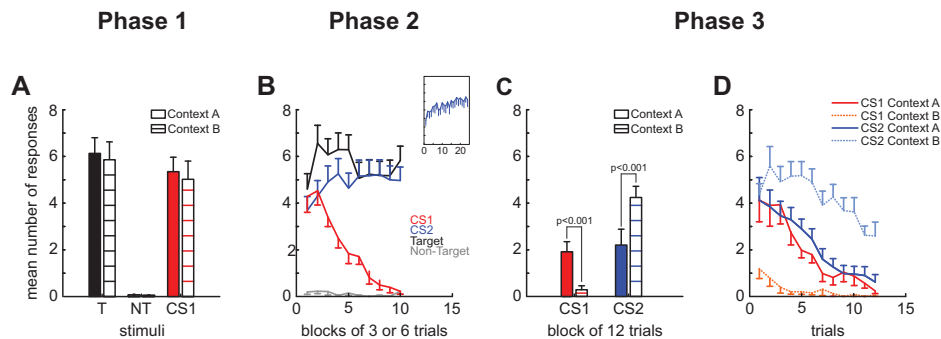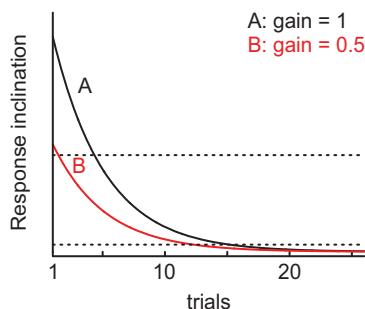


**Figure 2.** Absolute response counts (± SEM) during all three experimental phases. (*A*) Mean number of responses (± SEM) to the target and nontarget as well as to CS1-A and CS1-B in the last session of Phase 1. Solid bars show responding in context A while striped bars indicate responding in context B. (*B*) Mean number of responses (± SEM) during the second phase. Responding to different stimuli is color coded. Responding to the target (black) and non-target (gray) is shown in blocks of three trials, whereas responding to CS1 (red) and CS2 (blue) is averaged across six trials. We used different block sizes for display and analyses, because CS1 and CS2 were presented twice as often as T and NT in this phase. In addition, we pooled responses to CS1-A and CS1-B as well as CS2-A and CS2-B from this phase on. Inset shows mean response count (± SEM) to the CS2 for the first 24 trials. (*C*) Mean number of responses (± SEM) to CS1 and CS2 in both contexts, averaged across all 12 trials in Phase 3. For both stimuli, subjects responded more strongly in their respective acquisition context (A for CS1 and B for CS2), indicating renewal for CS1 and context specificity for CS2. (*D*) As in *C*, but response counts are shown for all trials individually. While there is a clear context effect for the CS1 at the beginning of the test session, responding to the CS2 differed between contexts from the second trial on.

earlier demonstration of contextual control of initial conditioning (Harris 2000).

How can the different temporal trajectories of context-specific responding for CS1 and CS2 in the present study be explained? Expanding on the ideas of the configural theory of Pearce (1994) in regard to unitary representations of context and stimulus and generalization based on similarity, we propose a model in which context modulates associative strength by attenuating its gain (see also Delamater and Westbrook 2014; Urcelay and Miller 2014 for similar ideas). In this model, associative strength is multiplied by a factor between 1 (context of learning) and 0 (maximally distinct context). This product defines the overall response inclination of the animals. The more different the context of testing is from that of learning, the lower the response inclination, given a fixed level of associative strength. Importantly, a reduction in response inclination (by either a context switch or extinction) might not directly be observable in conditioned responding: during acquisition in a single context, associative strength (and accordingly response inclination) may continue to build up when responding is already maximal. Conversely, associative strength may continue to decrease during extinction after overt responding has ceased (the classical learning–performance distinction; Soderstrom and Bjork 2015).

Figure 3 shows two extinction curves, starting at different levels of associative strength. The upper dashed horizontal line denotes the threshold below which changes in response inclination translate into observable changes in behavior. Curve A corresponds to a very strong association whose strength declines from the first trial on, although this decline would be observable in behavior only after four to five trials when inclination has decreased below the threshold (upper dotted line). This graph corresponds to what is observed for CS2 in its acquisition context: high levels of responding due to high associative strength without attenuation by a context switch (extinction in learning context). If responding to a CS is assessed in a different context (gain factor 0.5, curve B), the level of conditioned responding decreases after one or two nonreinforced presentations already. Our observation that level of responding to CS1 is on average lower than that to CS2 likely results from the two CSs' different reinforcement histories (extinction of CS1 but not CS2 in the previous session).

Even though it is widely accepted that extinction learning is more context specific than acquisition, quite a number of studies have shown contextual control of initial acquisition (Hall and Honey 1990; Rescorla 2008; Nelson 2009; Lucke et al. 2013). A recent series of experiments by Bouton et al. investigated the context specificity of instrumental learning and reported contextual control over nonextinguished behavior in terms of incomplete generalization from one context to another (Bouton et al. 2011, 2012, 2014). This challenges the claim that extinction is in general more context specific than acquisition. However, the authors report in addition AAB and ABC renewal effects, which may be interpreted as a sign of stronger contextual control of extinction memory. The authors attribute the difference concerning generalization of acquisition performance from one context to the other to the use of Pavlovian vs. instrumental responses, i.e., instrumental performance is proposed to be more context specific than Pavlovian responding (Bouton and Todd 2014). Our experiment adds to that body of research by providing data on Pavlovian conditioning in the appetitive domain, demonstrating that Pavlovian excitatory conditioned responses can be highly context specific under certain conditions.

In conclusion, we provide experimental data able to differentiate between influential learning theories based on their predictions concerning contextual control of behavior. While retrieval and elemental learning theories cannot account for the data, the configural theory is able to explain the present results. Finally, our model expands on configural ideas of unitary representations of context and CS by proposing that context influences responding by modulating the gain of associative strength.

## Acknowledgments

## References

Bouton ME. 2004. Context and behavioral processes in extinction. *Learn Mem* **11:** 485–494.

Bouton ME, King DA. 1983. Contextual control of the extinction of conditioned fear: tests for the associative value of the context. *J Exp Psychol Anim Behav Process* **9:** 248–265.

Bouton ME, Todd TP. 2014. A fundamental role for context in instrumental learning and extinction. *Behav Processes* **104:** 13–19.

Bouton ME, Westbrook RF, Corcoran KA, Maren S. 2006. Contextual and temporal modulation of extinction: behavioral and biological mechanisms. *Biol Psychiatry* **60:** 352–360.

Bouton ME, Todd TP, Vurbic D, Winterbauer NE. 2011. Renewal after the extinction of free operant behavior. *Learn Behav* **39:** 57–67.

Bouton ME, Winterbauer NE, Todd TP. 2012. Relapse processes after the extinction of instrumental learning: renewal, resurgence, and reacquisition. *Behav Processes* **90:** 130–141.

Bouton ME, Todd TP, León SP. 2014. Contextual control of discriminated operant behavior. *J Exp Psychol Anim Learn Cogn* **40:** 92–105.

Delamater AR, Westbrook RF. 2014. Psychological and neural mechanisms of experimental extinction: a selective review. *Neurobiol Learn Mem* **108:** 38–51.

Güntürkün O, Stüttgen MC, Manns M. 2014. Pigeons as a model species for cognitive neuroscience. *e-Neuroforum* **5:** 86–92.

Hall G, Honey RC. 1990. Context-specific conditioning in the conditioned-emotional-response procedure. *J Exp Psychol Anim Behav Process* **16:** 271–278.

Harris JA, Jones ML, Bailey GK, Westbrook RF. 2000. Contextual control over conditioned responding in an extinction paradigm. *J Exp Psychol Anim Behav Process* **26:** 174–185.

Lengersdorf D, Stüttgen MC, Uengoer M, Güntürkün O. 2014. Transient inactivation of the pigeon hippocampus or the nidopallium caudolaterale during extinction learning impairs extinction retrieval in an appetitive conditioning paradigm. *Behav Brain Res* **265:** 93–100.

Lengersdorf D, Marks D, Uengoer M, Stüttgen MC, Güntürkün O. 2015. Blocking NMDA-receptors in the pigeon's "prefrontal" caudal



**Figure 3.** Contextual gain model. Response inclination (RI) is plotted against trials in extinction. RI is defined as the product of associative strength (AS) and contextual gain (CG): RI = AS × CG. Color-coded curves correspond to differential modulatory influence by the context (*A* (black): gain = 1; *B* (red): gain = 0.5). Dashed horizontal lines denote the thresholds below and above which changes in the response inclination translate into effective changes in responding. At high levels, a reduction in response inclination might not directly be observable in responding (see main text for details).

nidopallium impairs appetitive extinction learning in a sign-tracking paradigm. *Front Behav Neurosci* **9:** 85.

Lovibond PF, Preston GC, Mackintosh NJ. 1984. Context specificity of conditioning, extinction, and latent inhibition. *J Exp Psychol Anim Behav Process* **10:** 360–375.

Lucke S, Lachnit H, Koenig S, Uengoer M. 2013. The informational value of contexts affects context-dependent learning. *Learn Behav* **41:** 285–297.

Nelson JB. 2009. Contextual control of first- and second-learned excitation and inhibition in equally ambiguous stimuli. *Learn Behav* **37:** 95–106.

Pearce JM. 1994. Similarity and discrimination: a selective review and a connectionist model. *Psychol Rev* **101:** 587–607.

Rescorla RA. 2008. Within-subject renewal in sign tracking. *Q J Exp Psychol* **61:** 1793–1802.

Rescorla RA, Wagner AR. 1972. A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In *Classical conditioning II: Current research and theory* (ed. Black AH, Prokasy WF), pp. 64–99. Appleton-Century-Crofts, New York.

Rosas JM, Callejas-Aguilera JE. 2006. Context switch effects on acquisition and extinction in human predictive learning. *J Exp Psychol Learn Mem Cogn* **32:** 461–474.

Rosas JM, Callejas-Aguilera JE, Ramos-Álvarez MM, Abad MJF. 2006. Revision of retrieval theory of forgetting: what does make information context specific? *Intern J Psych Psychol Therapy* **32:** 147–166.

Rosas JM, Todd TP, Bouton ME. 2013. Context change and associative learning. *Wiley Interdiscip Rev Cogn Sci* **4:** 237–244.

Soderstrom NC, Bjork RA. 2015. Learning versus performance: An integrative review. *Perspect Psychol Sci* **10:** 176–199.

Starosta S, Güntürkün O, Stüttgen MC. 2013. Stimulus-response-outcome coding in the pigeon nidopallium caudolaterale. *PLoS One* **8:** e57407.

Starosta S, Stüttgen MC, Güntürkün O. 2014. Recording single neurons' action potentials from freely moving pigeons across three stages of learning. *J Vis Exp* **88:** e51283.

Stüttgen MC, Yildiz A, Güntürkün O. 2011. Adaptive criterion setting in perceptual decision making. *J Exp Anal Behav* **96:** 155–176.

Stüttgen MC, Kasties N, Lengersdorf D, Starosta S, Güntürkün O, Jäkel F. 2013. Suboptimal criterion setting in a perceptual choice task with asymmetric reinforcement. *Behav Processes* **96:** 59–70.

Todd TP, Vurbic D, Bouton ME. 2014. Behavioral and neurobiological mechanisms of extinction in Pavlovian and instrumental learning. *Neurobiol Learn Mem* **108:** 52–64.

Üngör M, Lachnit H. 2006. Contextual control in discrimination reversal learning. *J Exp Psychol Anim Behav Process* **32:** 441–453.

Urcelay GP, Miller RR. 2014. The functions of contexts in associative learning. *Behav Processes* **104:** 2–12.